arXiv:1609.03436v1 [stat.ME] 12 Sep 2016

The Scalable Langevin Exact Algorithm: Bayesian Inference for Big Data

Murray Pollock[*], **Paul Fearnhead**[**],
Adam M. Johansen[†] **and Gareth O. Roberts**[‡]
e-mail: [*]m.pollock@warwick.ac.uk; [**]p.fearnhead@lancs.ac.uk
[†]a.m.johansen@warwick.ac.uk; [‡]gareth.o.roberts@warwick.ac.uk

Abstract: This paper introduces a class of Monte Carlo algorithms which are based upon simulating a Markov process whose quasi-stationary distribution coincides with the distribution of interest. This differs fundamentally from, say, current Markov chain Monte Carlo in which we simulate a Markov chain whose stationary distribution is the target. We show how to approximate distributions of interest by carefully combining sequential Monte Carlo methods with methodology for the exact simulation of diffusions. Our methodology is particularly promising in that it is applicable to the same class of problems as gradient based Markov chain Monte Carlo algorithms but entirely circumvents the need to conduct Metropolis-Hastings type accept/reject steps whilst retaining *exactness*: we have theoretical guarantees that we recover the correct limiting target distribution. Furthermore, this methodology is highly amenable to big data problems. By employing a modification to existing naïve subsampling techniques we can obtain an algorithm which is still exact but has *sub-linear* iterative cost as a function of data size.

Keywords and phrases: Control Variates, Importance sampling, Killed Brownian motion, Langevin diffusion, Markov chain Monte Carlo, Quasi-stationarity, Sequential Monte Carlo.

1. Introduction

Advances in methodology for the collection and storage of data have led to scientific challenges in a wide array of disciplines. This is particularly the case in Statistics as the complexity of appropriate statistical models and those we wish to fit often increases with data size. Many current state-of-the-art statistical methodologies have algorithmic cost that scale poorly with increasing volumes of data. As noted by [30], 'many statistical procedures either have unknown runtimes or runtimes that render the procedure unusable on large-scale data' and has resulted in a proliferation in the literature of methods '... which may provide no statistical guarantees and which in fact may have poor or even disastrous statistical properties'.

This is particularly keenly felt in computational and Bayesian statistics, in which the standard computational tools are Markov chain Monte Carlo (MCMC), Sequential Monte Carlo (SMC) and their many variants (see for example [49]). MCMC methods are *exact* in the (weak) sense that they construct Markov

chains which have the correct limiting distribution. Although MCMC methodology has had considerable success in being applied to a wide variety of substantive areas, they are not well-suited to this new era of 'big data' for the following reasons:

- Most MCMC algorithms require at least one likelihood evaluation per iteration. For example, Metropolis-Hastings algorithms require likelihood evaluations in order to perform the *accept/reject* step. Given a data set of size n, this will be at best an $\mathcal{O}(n)$ calculation. For algorithms which avoid accept/reject steps, such as Gibbs samplers, alternative $\mathcal{O}(n)$ calculations such as the evaluation of suitable sufficient statistics are required.
- A common approach in MCMC for missing data and latent variable models is data augmentation [54]. In the context of MCMC for large data sets this approach increases the dimensionality of the state space, and in some cases this increase will be $\mathcal{O}(n)$. As we demonstrate in Section 5.2 we are able to circumvent this issue.

Naturally, owing to the previous successes of MCMC, there has been considerable effort to modify existing methodology in order to address computational scalability. The success of such methods in reducing algorithmic cost has been mixed, and has invariably led to a compromise on exactness — such methodologies generally construct a stochastic process with limiting distribution which is (at least hopefully) close to the desired target distribution. Broadly speaking these methods can be divided into three classes of approach: 'Divide-and-conquer' methods; 'Exact Subsampling' methods; and, 'Approximate Subsampling' methods. Each of these approaches has its own strengths and weaknesses, and we will briefly review these in the following paragraphs.

Divide-and-conquer methods (for instance, [43, 59, 51, 40]) begin by splitting the data set into a large number of smaller data sets (which may or may not overlap). Inference is then conducted on these smaller data sets and resulting estimates are somehow combined. A clear advantage of such an approach is that inference on each small data set can be conducted independently, and in parallel, and so if one had access to a large cluster of computing cores then computational cost would be significantly reduced. Interaction between the cores is typically impracticable as the primary bottleneck is not computational speed, but latency. Therefore, from a parallelisation perspective this approach is perhaps the most appealing. The primary weakness of these methods is that the recombination of the separately conducted inferences is inexact. All current theory is asymptotic in the number of data points, n [43, 37]. For these asymptotic regimes the posterior will tend to a Gaussian distribution [29], and it is questionable whether divide-and-conquer methods offer an advantage over simple approaches such as a Laplace approximation to the posterior [4]. Most results on convergence rates (e.g. [52]) have rates that are of the $\mathcal{O}(m^{-1/2})$, where m is the number of datapoints in each sub-set. As such they are no stronger than convergence rates for analysing just a single batch. One exception is in [37], where convergence rates of $\mathcal{O}(n^{-1/2})$ are obtained, albeit under strong conditions. However these results

only relate to estimating marginal posterior distributions, rather than the full posterior.

Subsampling methods are designed so that each iteration requires access to only a subset of the data. Exact approaches in this vein typically require subsets of the data of random size at each iteration. One approach is to construct unbiased estimators of point-wise evaluations of the target density using subsets of the data, which could then be embedded within the pseudo-marginal MCMC framework recently developed by [1]. Unfortunately, the construction of such positive unbiased estimators is not possible in general [26]. An approach to MCMC in which the acceptance of a proposed move is determined by using a subset of the data is proposed by [35], but they choose the size of subset such thatwith high probability the correct decision is made. However, it was shown in [3] that such approaches, although cheaper, require subsets of the data of size $\mathcal{O}(n)$, and thus do not even reduce the computational complexity of standard MCMC.

Approximate subsampling methods have the advantage that they typically require fixed subsets of the data at every iteration, and are based on direct (although often difficult to justify) modifications of exact methods which themselves benefit from sound theoretical underpinning. In particular, one direction which has garnered considerable interest is to exploit the fact that the Langevin diffusion has known invariant distribution. If we denote our target density by π then the Langevin diffusion

$$\mathrm{d}\mathbf{X}_t = \frac{1}{2}\mathbf{\Lambda}\nabla \log \pi(\mathbf{X}_t)\,\mathrm{d}t + \mathbf{\Lambda}^{1/2}\,\mathrm{d}\mathbf{W}_t, \qquad \mathbf{X}_0 = \mathbf{x} \in \mathbb{R}^d, \;\; t \in [0, \infty), \tag{1}$$

in which $\mathbf{W}_t$ is d-dimensional Brownian motion and $\mathbf{\Lambda}$ is a prespecified preconditioning matrix, has π as its invariant distribution. So one could simply simulate this diffusion as a means to obtain Monte Carlo draws from π. Unfortunately, exact simulation from the Langevin diffusion is not possible for general π. Were it possible, even in the simplest Bayesian setting $\pi(\mathbf{x}) \propto \prod_{i=0}^{n} f_i(\mathbf{x})$ (where $f_0(\mathbf{x})$ is a prior and $\{f_i(\mathbf{x})\}_{i=1}^{n}$ are likelihood factors), evaluation of the drift of the diffusion would be an $\mathcal{O}(n)$ calculation.

One natural approach to big data problems suggested by [60] is to employ a stochastic gradient Euler-Maruyama discretisation scheme. As in a regular Euler-Maruyama discretisation scheme they suggest that the sample path is simulated iteratively over small time intervals of length $h > 0$, but where the drift component of the diffusion dynamics is estimated from a subset of the available data. So for a single iteration of their algorithm one would simulate

$$\mathbf{X}_{t+h} | (\mathbf{X}_t = \mathbf{x}) \sim \mathrm{N}\left(\mathbf{x} + h\tilde{\beta}(\mathbf{x}), h\mathbf{\Lambda}^{1/2}\right), \tag{2}$$

where $\tilde{\beta}$ is an unbiased estimator of the drift of (1) calculated at $\mathbf{X}_t = \mathbf{x}$ using a random subset of size $m \ll n$ of the available data. One would hope that the

approximation from using this sub-sampling estimate of the drift would be small if the variance of the estimator is small compared to the variance of the driving Brownian motion over a time-step, $h\mathbf{\Lambda}^{1/2}$. The attraction of this approach is that the construction of such a $\tilde{\beta}$ is trivial, requiring only the unbiased estimation of a sum. In particular, we might draw a random subset A of size m without replacement from $\{0, \ldots, n\}$, setting

$$\tilde{\beta}(\mathbf{x}) = \frac{n+1}{2m}\mathbf{\Lambda}\sum_{j\in A}\nabla \log f_j(\mathbf{x}) \ . \tag{3}$$

Such 'stochastic gradient' approximations are not exact: even if we were to use the entire data set at every iteration such an approach would, due to the discretisation error, not recover the correct target distribution [5]. A comprehensive theoretical study of the conditions and (fixed) step sizes required to recover a distribution which is close to the target distribution is provided by [22]. The most obvious way then to ensure the correct target distribution is obtained, is to embed the discretised dynamics as a proposal within a Markov chain Monte Carlo framework, which leads to the class of Metropolis-adjusted Langevin algorithms outlined in [5, 50]. However, due to the $\mathcal{O}(n)$ accept/reject step this MCMC embedding renders the approach impractical for large data sets.

Despite these theoretical concerns, using either (2)–(3) or some related approach [13, 38] has been shown to work well in practice. They can give accurate results and lead to speed-ups when compared to standard MCMC. However, for the approximation error from sub-sampling to be small we often require sub-sample sizes that are $\mathcal{O}(n)$. Furthermore, it can be hard to quantify the bias introduced by the both the subsampling and the time-discretisation of the Langevin dynamics. Some effort has been made to rectify the original naïve stochastic gradient scheme by using an increasingly fine discretisation with iteration [61, 58]. Although [22] show that this approach recovers (asymptotically) the correct target distribution, the trade-off is a computational cost which increases nonlinearly with diffusion time, and limited model class applicability (at least with any theoretical guarantees).

The approach to the problem of big data proposed below is a significant departure from the current literature. Rather than building our methodology upon the stationarity of appropriately constructed Markov chains, we develop a novel approach based on the *quasi-limiting* distribution of suitably constructed stochastically weighted diffusion processes. A *quasi-stationary distribution* for a Markov process X with respect to a Markov stopping time τ is the limit of the distribution of $X_t \mid \tau > t$ as $t \to \infty$ [16], and is completely unrelated to the popular area of Quasi-Monte Carlo. These *Quasi-Stationary Monte Carlo (QSMC) methods* which we develop can be used for a broad range of Bayesian problems (of a similar type to MCMC) and exhibit interesting and differing algorithmic properties. There are a number of different possible implementations of the theory which open interesting avenues for future research, in terms of branching

processes, by means of stochastic approximation methods, or (as outlined in this paper) SMC methods. One particularly interesting difference between our class of Monte Carlo algorithms and MCMC is that QSMC methods allow us to circumvent entirely the Metropolis-Hastings type accept/reject steps, while still retaining theoretical guarantees that the correct limiting target distribution is recovered. In the case of big data problems, this removes one of the fundamental $\mathcal{O}(n)$ bottlenecks in computation.

Our Quasi-Stationary Monte Carlo methods can be applied in big data contexts by using a novel subsampling approach. The resulting *Scalable Langevin Exact Algorithm (ScaLE)* not only has the correct limiting target distribution (and so in this sense is *exact*) but also has *sub-linear* iterative cost as a function of data size. As we show in Section 4, our approach to subsampling decreases the computational complexity of QSMC despite requiring subsampled data sets of fixed $\mathcal{O}(1)$ size. A key component of our approach is the use of control variates to reduce the variability of sub-sampling estimators of features of the posterior, but in addition exploits properties and simulation methods for Brownian motion. Although the use of control variates for MCMC in the context of big data and sub-sampling have been previously considered [4, 25], this is the first use of control variates within an exact method. We note that closely related control variate methods are being developed in concurrent work on MCMC (see [9]) and it seems the methodology may have quite wide applicability within the field of likelihood inference for big data.

The remainder of this paper is structured as follows. In the next section we outline the theory underpinning the methodology we develop, with Sections 2 and 2.1 introducing diffusions and the theory for diffusions typically exploited in Monte Carlo methodology, while Section 2.2 outlines the construction of our quasi-stationary process and the theory by which one can parametrise this process in order for it to recover a probability distribution of interest. Section 3 is focussed on developing Monte Carlo methodology based on the theory in Section 2. In particular, Section 3.1 outlines methodology for the exact simulation of diffusions based on path-space rejection sampling and their application to the simulation of our quasi-stationary process (without normalisation), along with an importance sampling variant. Section 3.2 provides a theoretical underpinning for embedding the importance sampling variant within a Sequential Monte Carlo (SMC) framework, along with algorithmic detail for our Quasi-Stationary Monte Carlo (QSMC) algorithm. In Section 4 we extend our QSMC method to deal with the particular problem of 'big data' and present the Scalable Langevin Exact Algorithm (ScaLE) which is an application of these methods to QSMC. We conclude in Section 5 by providing a number of examples to which we have applied our methodology. Note that for clarity of presentation we have suppressed much of the technical and algorithmic detail, but this can be found in the appendices: Appendix A examines the conditions on the theory of our quasi-stationarity of killed Brownian motion by considering heavy tailed densities; Appendix B provides conditions under which the path-space rejection

sampling construction of Section 3.1 can be employed; Appendix C details from basic random number generation the mechanism by which we simulate our process; Appendices D and E provide further technical and algorithmic detail on the path-space rejection sampler and killed Brownian motion algorithm respectively, which are introduced and discussed in Section 3.1; Appendix F provides detail of the rejection sampling variant of Section 3.1 within the SMC construction of Section 3.2; in Appendix G we provide details of the sub-sampling embedding within a rejection sampling framework; and finally, Appendix H provides more theoretical rigour to the application of SMC methodology we use.

2. Diffusions for Monte Carlo

A (d-dimensional) diffusion $\mathbf{X} : \mathbb{R} \to \mathbb{R}^d$ over an infinite-time horizon is a Markov process with almost surely continuous sample paths, which can be defined as the solution to the following stochastic differential equation (SDE),

$$\mathrm{d}\mathbf{X}_t = \beta(\mathbf{X}_t)\,\mathrm{d}t + \sigma(\mathbf{X}_t)\,\mathrm{d}\mathbf{W}_t \qquad \mathbf{X}_0 = \mathbf{x} \in \mathbb{R}^d, \;\; t \in [0, \infty), \tag{4}$$

where $\beta : \mathbb{R}^d \to \mathbb{R}^d$ denotes the (instantaneous) drift vector, $\sigma : \mathbb{R}^d \to \mathbb{R}_+^{d \times d}$ denotes the (instantaneous) volatility matrix and $\mathbf{W}_t$ is a standard (d-dimensional) Brownian motion. Regularity conditions are assumed to hold to ensure the existence of a unique non-explosive weak solution (see for instance [44]).

As discussed in Section 1, we are interested in exploiting the known theoretical properties of diffusions in order to sample from some distribution of interest. The Fokker–Planck partial differential equation describes the law of the diffusion (4). In particular, the evolution of the probability density for $\mathbf{X}_t$, denoted by $\nu(\mathbf{x}, t)$, can be described as follows,

$$\frac{\partial}{\partial t}\nu(\mathbf{x}, t) = -\sum_i \frac{\partial}{\partial \mathbf{x}_i}\left[\beta_i(\mathbf{x})\nu(\mathbf{x}, t)\right] + \frac{1}{2}\sum_{i,j}\frac{\partial^2}{\partial \mathbf{x}_i \partial \mathbf{x}_j}\left[\left[\sigma(\mathbf{x})\sigma^T(\mathbf{x})\right]_{i,j}\nu(\mathbf{x}, t)\right]. \tag{5}$$

We are interested in choosing the dynamics of the diffusion such that it has an invariant distribution that coincides with a distribution of interest. One can then, in principle, construct a scheme for generating samples suitable for approximating expectations with respect to that distribution. This is done by simply simulating a (finite horizon) sample trajectory from the measure induced on path-space by (4). Mirroring the use of stationary discrete time processes in an MCMC context, such an approach is motivated by the elementary fact that if $\mathbf{X}_t \sim \nu$, where ν is the invariant distribution, then $\mathbf{X}_{t+s}|(\mathbf{X}_t = \mathbf{x}_t) \sim \nu$ for every $s > 0$. As we are interested in the simulation of such a class of diffusions without any form of error, we further restrict our attention to the class of diffusions in which the volatility matrix is constant (i.e., $\sigma(\mathbf{x}) \equiv \mathbf{\Lambda}^{1/2}$). Setting $\nu(\mathbf{x}, t) = \nu(\mathbf{x})$ in the Fokker-Planck equation it is straightforward to show that

the following 'Langevin' diffusion has $\nu(\mathbf{x})$ as its invariant distribution,

$$\mathrm{d}\mathbf{X}_t = \frac{1}{2}\mathbf{\Lambda}\nabla \log \nu(\mathbf{X}_t)\,\mathrm{d}t + \mathbf{\Lambda}^{1/2}\,\mathrm{d}\mathbf{W}_t, \quad \mathbf{X}_0 := \mathbf{x} \sim \nu \in \mathbb{R}^d,\, t \in [0,\infty), \quad (6)$$

Furthermore [50] show that under certain regularity conditions the Langevin diffusion will converge at an exponential rate to the correct invariant distribution from any starting point $\mathbf{X}_0 := \mathbf{x} \in \mathbb{R}^d$.

2.1. Traditional Exploitation of Diffusions for Monte Carlo

Loosely speaking, in order for a computational statistician to draw from some target distribution π, one approach she could take (although not the only approach, as will become clear in the sequel) is to allow a diffusion to evolve with dynamics such that the invariant distribution ν coincides with π and then use its occupation measure as a proxy for π. This approach requires the simulation and storage of a continuous time object over an infinite time horizon, which is not feasible. We can, however, attempt to realise and store on a computer some finite-dimensional subset of the sample path noting that, as we can exploit the strong Markov property, we can simulate a single trajectory of arbitrary length by iteratively simulating and concatenating short sub-trajectories. We can then use the realised draws of the trajectory at the end of each sub-trajectory as a proxy for the occupation measure, and in turn realisations from π. As such, we can restrict our attention to the (iterative) simulation of finite-time horizon diffusions of the following form,

$$\mathrm{d}\mathbf{X}_t = \beta(\mathbf{X}_t)\,\mathrm{d}t + \mathbf{\Lambda}^{1/2}\,\mathrm{d}\mathbf{W}_t, \quad \mathbf{X}_0 = \mathbf{x} \in \mathbb{R}^d,\, t \in [0,T], \quad (7)$$

where $\beta(\mathbf{x}) := \frac{1}{2}\mathbf{\Lambda}\nabla \log \nu(\mathbf{x}) :\ \mathbb{R}^d \to \mathbb{R}^d$ denotes the (instantaneous) drift vector. In the following let $\mathbb{Q}^{\mathbf{x}}$ denote the probability distribution induced on path-space by (7), and $\mathbb{Q}^{\mathbf{x},\mathbf{y}}$ denote $\mathbb{Q}^{\mathbf{x}}$ conditioned on $\mathbf{X}_T = \mathbf{y} \in \mathbb{R}$.

The methodological approach outlined above requires us to be able to simulate from the transition density of (7), which we denote by $p_T(\mathbf{x},\mathbf{y}) := \mathbb{P}_{\mathbb{Q}^{\mathbf{x}}}(\mathbf{X}_T \in \mathrm{d}\mathbf{y} \,|\, \mathbf{X}_0 = \mathbf{x})/\,\mathrm{d}\mathbf{y}$. Under conditions given in Appendix B, an explicit (although not tractable) form for the transition density can be found as detailed in Corollary 2. Allowing $\mathbb{W}^{\mathbf{x}}$, $\mathbb{W}^{\mathbf{x},\mathbf{y}}$ to denote the law of a driftless version of (7) and its conditioned counterpart, defining $w_T(\mathbf{x},\mathbf{y}) := \mathbb{P}_{\mathbb{W}^{\mathbf{x}}}(\mathbf{X}_T \in \mathrm{d}\mathbf{y} \,|\, \mathbf{X}_0 = \mathbf{x})/\,\mathrm{d}\mathbf{y}$ and $\phi(\mathbf{u}) := (||\mathbf{\Lambda}^{-1/2}\beta(\mathbf{u})||^2 + \operatorname{div}\mathbf{\Lambda}^{-1/2}\beta(\mathbf{u}))/2$, where $||\cdot||$ denotes the Euclidean norm, we can write

$$p_T(\mathbf{x},\mathbf{y}) = w_T(\mathbf{x},\mathbf{y}) \cdot \mathbb{E}_{\mathbb{W}^{\mathbf{x},\mathbf{y}}}\left[\frac{\mathrm{d}\mathbb{Q}^{\mathbf{x}}}{\mathrm{d}\mathbb{W}^{\mathbf{x}}}(\mathbf{X})\right] \quad (8)$$

$$\propto w_T(\mathbf{x},\mathbf{y}) \cdot \nu^{1/2}(\mathbf{y}) \cdot \mathbb{E}_{\mathbb{W}^{\mathbf{x},\mathbf{y}}}\left[\exp\left\{-\int_0^T \phi(\mathbf{X}_s)\,\mathrm{d}s\right\}\right]. \quad (9)$$

Furthermore, we have the classic stationarity result

$$\lim_{T\to\infty} p_T(\mathbf{x},\cdot) \to \nu(\cdot), \qquad \forall \mathbf{x}\in\mathbb{R}^d, \tag{10}$$

in a sense that will be made precise later, under suitable regularity conditions (see for example the discussion in [23]).

Now, adopting the approach outlined above in which the dynamics of the Langevin diffusion are chosen such that the invariant distribution ν coincides with some target distribution π, we require the ability to simulate from (9). One method to simulate from (9) is to deploy a numerical discretisation scheme [27, 33, 45] (a so-called 'Unadjusted Langevin Algorithm' (ULA)). However in all but a small number of trivial cases such discretisation schemes will not provide an exact draw from the transition density, and more importantly will not recover the correct invariant distribution (it may not even possess any invariant distribution [5]). It has been shown that the ULA can behave very badly [50].

An alternative to a discretisation scheme to simulate from the transition density is to deploy a rejection sampling scheme (the so called *Path-Space Rejection Sampler* (PRS) for diffusions [8, 7, 46, 12, 48]). Such a rejection sampling scheme would involve simulating from the alternative proposal transition density $q(\mathbf{y}) \propto w_T(\mathbf{x},\mathbf{y})\cdot\nu^{1/2}(\mathbf{y})$, and accepting or rejecting the proposal with probability proportional to $\mathbb{E}_{\mathbb{W}^{\mathbf{x},\mathbf{y}}}[\exp\{-\int_0^T \phi(\mathbf{X}_s)\,\mathrm{d}s\}]$. As detailed in Appendix B all the conditions necessary to implement a PRS are met for the class of Langevin diffusions we consider (6), but the implementation of such schemes would require the construction of a proposal density which incorporates the invariant distribution ν. If this were readily feasible, there would be no need to use a ν-invariant diffusion at all, and so the implementation of such a scheme is not practical.

This paper provides alternative methodology to access π by means of simulating normalised killed Brownian motion, entirely avoiding the issues outlined above associated with discretisation and rejection-based schemes.

To begin briefly outlining our approach (further details can be found in Section 2.2), we let $\{\mathbb{K}_T^{\mathbf{x}},\ T\geq 0\}$ be a collection of probability measures where $\mathbb{K}_T^{\mathbf{x}}$ describes a probability law on $C[0,T]$ such that the following Radon-Nikodým derivative exists,

$$\frac{\mathrm{d}\mathbb{K}_T^{\mathbf{x}}}{\mathrm{d}\mathbb{W}_T^{\mathbf{x}}}(\mathbf{X}) \propto \exp\left\{-\int_0^T \phi(\mathbf{X}_s)\,\mathrm{d}s\right\}. \tag{11}$$

Letting the marginal density of $\mathbb{K}_T^{\mathbf{x}}$ at time T be denoted by $k_T(\mathbf{x}, \mathbf{y})$, we establish that,

$$k_T(\mathbf{x}, \mathbf{y}) \propto w_T(\mathbf{x}, \mathbf{y}) \cdot \underbrace{\mathbb{E}_{\mathbb{W}_T^{\mathbf{x},\mathbf{y}}}\left[\frac{\mathrm{d}\mathbb{K}_T^{\mathbf{x}}}{\mathrm{d}\mathbb{W}_T^{\mathbf{x}}}(\mathbf{X})\right]}_{=:P(\mathbf{X})}, \tag{12}$$

which follows by noting that if, say, ϕ is everywhere non-negative, then $\mathbb{K}_T^{\mathbf{x}}$ is the law of killed Brownian motion conditioned to survive until time T with instantaneous killing rate $\phi(\mathbf{X}_s)$. Intuitively, recalling (9) and (10), we may naturally posit that

$$\lim_{T\to\infty} k_T(\mathbf{x}, \mathbf{y}) \propto p_T(\mathbf{x}, \mathbf{y}) \cdot \nu^{-1/2}(\mathbf{y}) \to \nu^{1/2}(\mathbf{y}), \tag{13}$$

and so if $\nu^{1/2}$ is itself integrable we may expect that $k_T(\mathbf{x}, \cdot)$ converges to a probability density proportional to $\nu^{1/2}$. Indeed, this is established in the following theorem,

Theorem 1 (Quasi-Stationarity of Killed Brownian motion)**.** *Suppose that $\nu^{1/2} \in L^1$, is bounded, and there exists some $\gamma \geq 1/2$ such that*

$$\liminf_{\mathbf{x}\to\infty}\left(\frac{\Delta\nu(\mathbf{x})}{\nu^{\gamma+1/2}(\mathbf{x})} - \frac{\gamma\|\nabla\nu(\mathbf{x})\|^2}{\nu^{\gamma+3/2}(\mathbf{x})}\right) > 0 \tag{14}$$

where Δ represents the Laplacian:

$$\Delta f(\mathbf{x}) = \sum_{i=1}^{d} \frac{\partial^2 f}{\partial x_i^2},$$

then we have the following convergence in L^1:

$$\lim_{T\to\infty} k_T(\mathbf{x}, \cdot) = \frac{\nu^{1/2}(\cdot)}{\int_{\mathbb{R}}^{d} \nu^{1/2}(\mathbf{z})d\mathbf{z}}. \tag{15}$$

Proof. See Section 2.2. □

Consequently, a computational statistician who wishes to draw from a target distribution π has a novel approach to explore. She can choose the dynamics of (6) such that the invariant distribution $\nu \equiv \pi^2$, and generate a trajectory of killed Brownian motion from $\mathbb{K}_T^{\mathbf{x}}$. Then by Theorem 1 we have (noting as required we have $\nu^{1/2} \in L^1$ as $\nu^{1/2} \equiv \pi$), as $T \to \infty$:

$$k_T(\mathbf{x}, \cdot) \xrightarrow{L^1} \pi(\cdot), \qquad \forall \mathbf{x} \in \mathbb{R}^d. \tag{16}$$

To generate a sample trajectory from $\mathbb{K}_T^{\mathbf{x}}$, we need only sample Brownian motion killed instantaneously at a state-dependent rate $\phi(\mathbf{X}_s)$, and this can be done by thinning a dominating Poisson process. This provides an unexplored

avenue for sampling (asymptotically) from a target distribution π without the need to resort to an unknown approximation.

In the remainder of this section we provide a proof and discussion of Theorem 1. We return in Section 3 to provide detail on how to exploit practically Theorem 1 and (16).

2.2. A Novel Quasi-Stationary Approach for Monte Carlo

In Section 2.1 we obtained an interesting limit result for the distribution of killed Brownian motion at a given time conditioned on not being killed by that time — a so called *quasi-limiting* distribution. The literature on quasi-limiting distributions and quasi-stationarity is vast (see for example [16] for an entry into the area) and contains many unexpected subtleties which distinguish it from the theory of stationarity for Markov chains and processes.

Here we present a proof of Theorem 1.

Proof (Theorem 1). The diffusion in (6) has generator given by

$$\mathfrak{A}f(\mathbf{x}) = \frac{1}{2}\Delta f(\mathbf{x}) + \frac{1}{2}\nabla \log \nu(\mathbf{x}) \cdot \nabla f(\mathbf{x}).$$

As ν is bounded, we assume without loss of generality that its upper bound is 1. Our proof shall proceed by checking the conditions of Corollary 6 of [23], which establishes the result: we need to check that the following are satisfied:

1. For all $\delta > 0$, the discrete time chain $\{X_{n\delta}, n = 0, 1, 2, \ldots\}$ is irreducible;
2. All closed bounded sets are petite;
3. We can find a drift function $V(\mathbf{x}) = \nu(\mathbf{x})^{-\gamma}$ for some $\gamma > 0$, that satisfies the condition

$$\mathfrak{A}V^{\eta}(\mathbf{x}) \leq -c_{\eta}V(\mathbf{x})^{\eta-\alpha} \tag{17}$$

for $\mathbf{x}$ outside some bounded set, for each $\eta \in [\alpha, 1]$ with associated positive constant c_{η} and where $\alpha = 1 - (2\gamma)^{-1}$.

The first condition holds for any regular diffusion since the diffusion possesses positive continuous transition densities over time intervals $t > 0$; and positivity and continuity of the density also implies the second condition. For the final condition we require that

$$\limsup_{|\mathbf{x}|\to\infty} \frac{\mathfrak{A}V^{\eta}(\mathbf{x})}{V^{\eta-\alpha}(\mathbf{x})} < 0 \ . \tag{18}$$

Now by direct calculation

$$\mathfrak{A}V^{\eta}(\mathbf{x}) = \frac{\eta\gamma}{2}\nu(\mathbf{x})^{-\gamma\eta-2}\left[\eta\gamma\|\nabla\nu(\mathbf{x})\|^2 - \nu(\mathbf{x})\Delta\nu(\mathbf{x})\right] \tag{19}$$

so that

$$\frac{\mathfrak{A}V^{\eta}(\mathbf{x})}{V(\mathbf{x})^{\eta-\alpha}} = \frac{\eta\gamma\nu(\mathbf{x})^{-3/2-\gamma}}{2}\left[\eta\gamma\|\nabla\nu(\mathbf{x})\|^2 - \nu(\mathbf{x})\Delta\nu(\mathbf{x})\right]. \tag{20}$$

Therefore (18) will hold whenever (14) is true since we have the constraint that $\eta \leq 1$ and $\|\nabla\nu(\mathbf{x})\|^2$ is clearly non-negative. As such the result holds as required. □

Note that the condition in (14) is essentially a condition on the tail of ν. This will hold even for heavy-tailed distributions, and we show this is the case for a class of 1-dimension target densities in Appendix A.

3. Exact Monte Carlo for Diffusions

In this section we develop Monte Carlo methodology for the practical exploitation of the theory developed in Section 2. We begin in Section 3.1 by outlining methodology for the exact simulation of diffusions based on path-space rejection sampling, with particular focus on their application to our desired quasi-stationary process. In Section 3.1 we focus on an importance sampling construction which requires a number of trajectories in order to conduct normalisation. The drawback of this approach is the variance of the importance sampling weights increases with diffusion time. We address this issue in Section 3.2 by embedding the methodology developed in Section 3.1 within a Sequential Monte Carlo (SMC) framework, providing a theoretical underpinning for doing so. In Section 3.2 we also provide algorithmic detail for the resulting *Quasi-Stationary Monte Carlo (QSMC)* algorithm.

We omit at this stage any discussion of big data, returning to this context in Section 4.

3.1. Path-space Rejection Samplers

Path-space Rejection Samplers (PRS) is the name we give to a class of rejection samplers operating on diffusion path-space over a finite time horizon, in which appropriately weighted finite dimensional subsets of sample paths are drawn from some target measure (in our case $\mathbb{K}_T^{\mathbf{x}}$), by means of simulating from a tractable equivalent measure with respect to which the target has bounded Radon-Nikodým derivative. As established in Theorem 1, one such equivalent measure is provided by Brownian motion with an appropriate pre-conditioning matrix $\mathbf{\Lambda}$, denoted by $\mathbb{W}^{\mathbf{x}}$. In this paper for simplicity we restrict our attention to diagonal pre-conditioning matrices, $\mathbf{\Lambda}$, but note there are no theoretical or methodological limitations which would prevent us from considering other non-diagonal matrices. $\mathbb{W}^{\mathbf{x}}$ is a convenient measure as finite dimensional subsets of sample paths can be drawn without error (by means of simulating scaled standard d-dimensional Brownian motion, see [46, §2.8]). In the interest of clarity

we defer, wherever possible, the technical conditions of PRS to Appendix B.

We first consider the simplest path-space rejection sampling setting (introduced by [8]). Provided that there exists $M < \infty$ such that $\frac{\mathrm{d}\mathbb{K}_T^{\mathbf{x}}}{\mathrm{d}\mathbb{W}_T^{\mathbf{x}}}(\mathbf{X}) \leq M$, we can proceed as in standard rejection sampling and draw $\mathbf{X} \sim \mathbb{W}^{\mathbf{x}}$ and accept the sample path ($I = 1$) with probability $P(\mathbf{X}) := \frac{1}{M}\frac{\mathrm{d}\mathbb{K}_T^{\mathbf{x}}}{\mathrm{d}\mathbb{W}_T^{\mathbf{x}}}(\mathbf{X}) \in [0,1]$, then $(\mathbf{X}|I=1) \sim \mathbb{K}_T^{\mathbf{x}}$. Under Condition 4 of Appendix B, such a bound can be found and

$$P(\mathbf{X}) = e^{\Phi T} \cdot \exp\left\{-\int_0^T \phi(\mathbf{X}_s)\,\mathrm{d}s\right\} \in [0,1]. \tag{21}$$

Direct simulation of an event of probability $P(\mathbf{X})$ is not possible as evaluation of the integral in (21) requires the realisation of an entire (infinite dimensional) sample path. However, as noted in [8, 6, 7, 46, 12, 48], an indirect and unbiased approach to simulate $P(\mathbf{X})$, which can be implemented using only a finite dimensional realisation of the sample path is possible.

This approach to simulate $P(\mathbf{X})$ requires upper and lower bounds for $\phi(\mathbf{X}_{[0,T]})$ (to ensure that it can be computed in finite expected time). Such a condition on $\phi(\mathbf{X}_{[0,T]})$ is not restrictive as it is possible to simulate $\mathbf{X} \sim \mathbb{W}^{\mathbf{x}}$ in conjunction with a sequence of stopping times $\{\tau_0 := 0, \tau_1, \tau_2, \ldots\}$, such that for the interval between any two successive stopping times the sample path is constrained to a hypercuboid, $\mathbf{X}_{[\tau_{i-1},\tau_i]} \in [\ell_{\mathbf{X}}^{(i)}, \upsilon_{\mathbf{X}}^{(i)}]$ where $\ell_{\mathbf{X}}^{(i)}, \upsilon_{\mathbf{X}}^{(i)} \in \mathbb{R}^d$ and for any $a = (a_1, \ldots, a_d), b = (b_1, \ldots, b_d) \in \mathbb{R}^d$, $[a,b] := \otimes_{i=1}^d [a_i, b_i]$. As a consequence of Result 2, such a *path-space layer* provides upper and lower bounds on $\phi(\mathbf{X})$ over any finite time interval (which will vary depending on the interval being considered).

The construction and simulation of the path-space layer is provided in detail in Appendix C. The user chooses the extent of the state space inscribed by the hypercuboid, $[\ell_{\mathbf{X}}^{(i)}, \upsilon_{\mathbf{X}}^{(i)}]$, in each time-interval. The algorithm then simulates the time and location, $(\tau_i, \mathbf{X}_{\tau_i})$, at which the process first leaves this hypercuboid. Furthermore, the construction detailed in Appendix C allows one to simulate the sample path at any other desired intermediate point. As such we further denote by $R_{\mathbf{X}}^{(i)}$ the path-space information and by $U_{\mathbf{X}}^{(i)} \in \mathbb{R}$ and $L_{\mathbf{X}}^{(i)} \in \mathbb{R}$ as the upper and lower bounds for $\phi(\mathbf{X}_{[\tau_{i-1},\tau_i]})$ (i.e. for the i^{th} stopping time interval, $[\tau_{i-1}, \tau_i] \subseteq [0,T]$). Conditional on the simulated layer ($R := R(\mathbf{X}) \sim \mathcal{R}$, comprised of $n_R := \inf\{n : \tau_n \geq T\}$ segments) we can represent our acceptance

probability for the draw $\mathbf{X} \sim \mathbb{W}_T^{\mathbf{x}}$ as follows

$$P(\mathbf{X}) = \underbrace{\exp\left\{\Phi T - \sum_{i=1}^{n_R} L_{\mathbf{X}}^{(i)} \cdot [(\tau_i \wedge T) - \tau_{i-1}]\right\}}_{=:P^{(1)}(\mathbf{X}) \in [0,1]} \cdot \prod_{i=1}^{n_R} \left[\underbrace{\exp\left\{ -\int_{\tau_{i-1}}^{\tau_i \wedge T} \left(\phi(\mathbf{X}_s) - L_{\mathbf{X}}^{(i)}\right) \mathrm{d}s \right\}}_{=:P^{(2,i)}(\mathbf{X})} \right] \tag{22}$$

$$= \prod_{i=1}^{n_R} \left[\underbrace{\exp\left\{(\Phi - L_{\mathbf{X}}^{(i)}) \cdot [(\tau_i \wedge T) - \tau_{i-1}]\right\}}_{=:P^{(1,i)}(\mathbf{X}) \in [0,1]} \cdot \underbrace{\exp\left\{ -\int_{\tau_{i-1}}^{\tau_i \wedge T} \left(\phi(\mathbf{X}_s) - L_{\mathbf{X}}^{(i)}\right) \mathrm{d}s \right\}}_{=:P^{(2,i)}(\mathbf{X})} \right]. \tag{23}$$

The path-space layer ensures $\phi(\mathbf{X}_{[\tau_{i-1},\tau_i]}) \in [L_{\mathbf{X}}^{(i)}, U_{\mathbf{X}}^{(i)}]$, and so $\forall i$

$$P^{(2,i)}(\mathbf{X}) \in \left[e^{-(U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)})T}, 1 \right] \subseteq (0, 1]. \tag{24}$$

From (22) note that the simulation of an event of probability $P(\mathbf{X})$ can be conducted by first simulating an event of probability $P^{(1)}(\mathbf{X})$ (which is straightforward to simulate as it is directly determined using the path-space layer simulated), and then n_R events of probability $P^{(2,i)}(\mathbf{X})$ (which are functions of the infinite-dimensional proposal sample path). Acceptance of $P(\mathbf{X})$ is equivalent to accepting all events in this extended construction. We can simulate an event of probability $P^{(2,i)}(\mathbf{X})$ indirectly and using a finite dimensional realisation of the sample path by means of the following Theorem 2.

Theorem 2 (Segment Acceptance Probability)**.** *Let $\mathbb{U}_R^{(i)}$ be the distribution of the tuple $\{\xi_1^{(i)}, \xi_2^{(i)}, \dots\} \cap [\tau_{i-1}, (\tau_i \wedge T)]$, where $\xi_j^{(i)} = \xi_{j-1}^{(i)} + \zeta_j^{(i)}$, $\zeta_j^{(i)} \sim \mathrm{Exp}(U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)})$ and $\xi_0^{(i)} := \tau_{i-1}$. Denoting by κ_i the cardinality of this tuple, we have,*

$$P^{(2,i)}(\mathbf{X}) = \mathbb{E}_{\mathbb{U}_R^{(i)}} \left[\underbrace{\prod_{j=1}^{\kappa_i} \left(\frac{U_{\mathbf{X}}^{(i)} - \phi\left(\mathbf{X}_{\xi_j^{(i)}}\right)}{U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)}} \right)}_{=:P^{(2,i,j)}(\mathbf{X})} \,\middle|\, \mathbf{X} \right]. \tag{25}$$

Proof. Follows directly from [48, §3.1], noting the law induced by the uniform distribution of a heterogenous Poisson process with intensity which changes at the time of each event so that it is, informally, $U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)}$ on the interval $[\tau_{i-1}, (\tau_i \wedge T)]$ has an equivalent distribution to the procedure above (see for instance [46, 32]). □

Theorem 2 allows us to simulate a Bernoulli random variable of success probability $P^{(2,i)}(\mathbf{X})$, by sampling the sample path at a finite number of exponentially distributed times, and simulating events of probability $P^{(2,i,j)}(\mathbf{X})$ (all of which need to be accepted for success). As such we now have a path-space rejection sampler for $\mathbb{K}_T^{\mathbf{x}}$. Over a pre-specified time horizon, then (22) and Theorem 2

together allow us to determine acceptance of $\mathbf{X} \sim \mathbb{W}_T^{\mathbf{x}}$, but only using a finite dimensional path skeleton. For clarity of presentation we omit the technical and algorithmic details of our PRS for $\mathbb{K}_T^{\mathbf{x}}$, which can instead be found in Appendix D and Algorithm 7.

Recall that we are ultimately interested in the transition density $k_T(\mathbf{x}, \mathbf{y})$ when $T \to \infty$. As such, we may more naturally be interested in identifying the time and location at which the Brownian motion sample path proposal is killed (denoted by $(\bar{\tau}, \mathbf{X}_{\bar{\tau}})$), as opposed to determining over a pre-specified time interval whether or not killing has occurred (as in the PRS for $\mathbb{K}_T^{\mathbf{x}}$ outlined above). Note that killing will almost-surely occur in finite time (unless $\Phi = U_{\mathbf{X}}$). As the proposal process is strongly Markov, by appealing to (23) the following is an informal description of a valid (although not unique) weight process in the interval $[\tau_{i-1}, \tau_i]$,

$$\mathrm{d}w_t = \begin{cases} -w_{t-} \cdot (L_{\mathbf{X}}^{(i)} - \Phi) \cdot \mathrm{d}t & \text{w.p. } (1 - \mathrm{d}t) \cdot (U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)}), \\ -w_{t-} \cdot \dfrac{\phi(\mathbf{X}_t) - L_{\mathbf{X}}^{(i)}}{U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)}} & \text{w.p. } \mathrm{d}t \cdot (U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)}). \end{cases} \tag{26}$$

The first line of (26) can be interpreted as a constant weight degradation in an increment of time arising from $P^{(1,i)}(\mathbf{X})$ of (23), and the second line of (26) the weight degradation if the increment coincides with an event time $\xi_\cdot^{(i)}$ as in Theorem 2.

Consequently, identifying the pair $(\bar{\tau}, \mathbf{X}_{\bar{\tau}})$ is trivial once one identifies the killing time (where $u \sim \mathrm{U}[0,1]$),

$$\bar{\tau} := \inf\{t \in [0, \infty) : 1 - w_t(\mathbf{X}_{[0,\infty)}) \leq u\}. \tag{27}$$

Fortuitously the simulation techniques available to us for our Brownian motion proposal constrained to a path-space layer (as detailed in Appendix C), is sufficiently tractable to allow us to simulate $(\bar{\tau}, \mathbf{X}_{\bar{\tau}})$ exactly. We term this algorithm *Killed Brownian Motion (KBM)*, full details of which can be found in Appendix E and Algorithm 8.

The Monte Carlo algorithm we focus on in this paper is an importance sampling variant of the KBM algorithm described above, in which Brownian motion with path-space information is simulated and weighted continuously in time, with evaluations of the trajectory occurring at exponentially distributed times as given in Theorem 2. If $L_{\mathbf{X},j} := \{L_{\mathbf{X}}^{(i)} : \xi_j \in [\tau_{i-1}, \tau_i]\}$ and $U_{\mathbf{X},j} := \{U_{\mathbf{X}}^{(i)} : \xi_j \in [\tau_{i-1}, \tau_i]\}$), then the weight of the trajectory at time $t \in [0, \infty)$ is simply

$$w_t(\mathbf{X}) = \exp\{\Phi t\} \cdot \underbrace{\exp\left\{ -\sum_{i : t \in [0, \tau_i]} L_{\mathbf{X}}^{(i)} \cdot [(\tau_i \wedge t) - \tau_{i-1}] \right\} \cdot \prod_{j : \xi_j \leq t} \left(\frac{U_{\mathbf{X},j} - \phi(\mathbf{X}_{\xi_j})}{U_{\mathbf{X},j} - L_{\mathbf{X},j}} \right)}_{=: w_t^*(\mathbf{X})}. \tag{28}$$

In practice we only need to calculate this weight up to a common constant of proportionality, so we calculate $w_t^*(\mathbf{X})$. We term this approach the Importance Sampling Killed Brownian Motion (IS-KBM) algorithm, which we present in Algorithm 1 and schematically in Figure 1. The advantage of this approach over KBM (presented in Appendix E) is that we can recover for any $t \in [0, \infty)$ a weighted trajectory, but to do so we need to halt the algorithm. If we require the trajectory at time t, this can be done by simply setting $\xi_j := \xi_j \wedge t$ after Algorithm 1 Step 4, halting upon simulation of the trajectory at time t, and returning $\mathbf{X}_t$ along with the (un-normalised) importance weight w_t^*.

Repeating this procedure several times allows us to use self-normalised importance sampling to approximate expectations with respect to the law of the killed process conditioned to remain alive until any desired time.

Algorithm 1 Importance Sampling Killed Brownian Motion (IS-KBM) Algorithm

1. Initialise: Input initial value $\mathbf{X}_0$. Set $i = 1$, $j = 0$, $\tau_0 = 0$, $w_0^* = 1$.
2. R: Simulate layer information $R_{\mathbf{X}}^{(i)} \sim \mathcal{R}$ as per Appendix C, obtaining τ_i, $L_{\mathbf{X}}^{(i)}$, $U_{\mathbf{X}}^{(i)}$.
3. ζ: Simulate $\zeta \sim \text{Exp}(U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)})$.
4. ξ_j: Set $j = j + 1$ and $\xi_j = (\xi_{j-1} + \zeta) \wedge \tau_i$.
5. $w_{\xi_j}^*$: Set $w_{\xi_j}^* = w_{\xi_{j-1}}^* \cdot \exp\{-L_{\mathbf{X}}^{(i)}[\xi_j - \xi_{j-1}]\}$.
6. $\mathbf{X}_{\xi_j}$: Simulate $\mathbf{X}_{\xi_j} \sim \text{MVN}(\mathbf{X}_{\xi_{j-1}}, (\xi_j - \xi_{j-1})\mathbf{\Lambda}) | R_{\mathbf{X}}^{(i)}$.
7. τ_i: If $\xi_j = \tau_i$, set $i = i + 1$, and return to Step 2. Else set $w_{\xi_j}^* = w_{\xi_j}^* \cdot (U_{\mathbf{X}}^{(i)} - \phi(\mathbf{X}_{\xi_j}))/(U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)})$ and return to Step 3 .

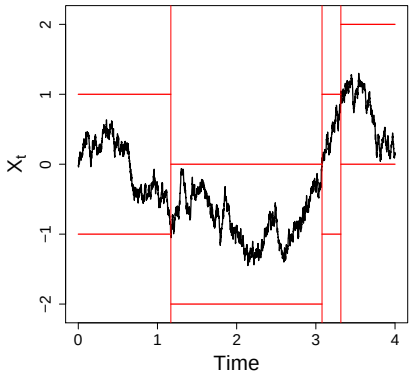

(a) Layers simulated for an illustrative (unrealised) Brownian motion sample path. Red horizontal lines mark the bounds of the layer for a time segment, red vertical lines mark stopping times in which there is a change in layer.

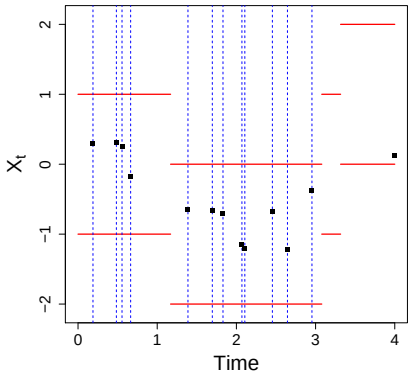

(b) Brownian motion sample path simulated at the finite collection of event times (black dots) as given by exogenous randomness, and path-space layers (red lines), that are required within IS-KBM.

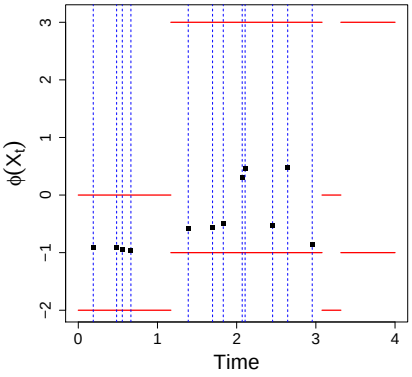

(c) Evaluation of ϕ at the event times of the sample path given by (b).

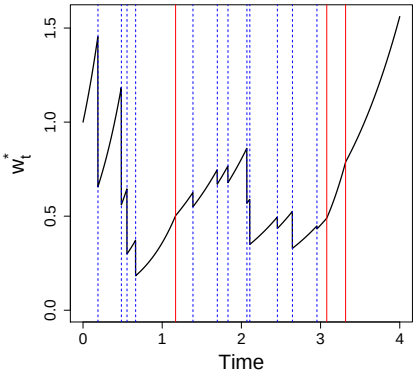

(d) Un-normalised importance weight process of sample path, comprising exponential growth between event times and discrete jumps at event times.

FIGURE 1. *Schematic of the Importance Sampling Killed Brownian Motion (IS-KBM) Algorithm in 1 dimension, simulated over an interval of length $T = 4$. To summarise, we would begin by simulating a set of path-space layers over the desired interval (as shown in (a)), realising the sample path at a collection of event times given by exogenous randomness (as shown in (b)), evaluating ϕ at those event times (as shown in (c)), and then recovering the (un-normalised) weight process which has exponential growth between weight times (with exponent $-L_{\mathbf{X}}^{(i)}$ as given by the associated path-space layer for ϕ) and discrete jumps at an event time ξ of size $(\phi(X_\xi) - L_{\mathbf{X}}^{(i)})/(U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)})$ (as shown in (d)).*

3.2. Normalisation through Sequential Monte Carlo

Although the IS-KBM framework we have presented thus far is conceptually appealing, in order to normalise we require a number of trajectories, and the variance of the importance sampling weights of these trajectories will increase with time. We present one possible strategy for addressing this difficulty in this section.

Sequential Monte Carlo (SMC) methods combine importance sampling and resampling techniques in order to approximate a sequence of distributions of interest. First used by [24] in the context of target tracking, such methods apply naturally when one is interested in approximating particular time marginals of a sequence of distributions defined upon an increasing state space (see, for example, [21, Section 24.3] for a general presentation). They propagate an ensemble of N weighted samples over time, attaching importance weights to each sample and using resampling to replicate samples with large weights and to eliminate those with small weights according to a stochastic rule which ensures that the sample remains properly weighted for the actual target distribution. They are particularly applicable in the current setting as the resampling step allows computational effort to be focussed on trajectories which have survived. Loosely speaking, one can think of resampling as a mechanism in which samples with low survival probability are eliminated and those with a high probability are allowed to branch.

We apply such an algorithm here in order to provide, iteratively, a population of samples (or *particles*) suitable for approximating the law of the killed process conditioned to survive until the current time. This can be done for both the rejection sampling (KBM) and importance sampling (IS-KBM) versions of the algorithm presented in Section 3.1. For clarity we present details of the KBM embedding in SMC in Appendix F, and instead consider here the importance sampling version as it more closely aligns with typical SMC constructions.

Considering the embedding of IS-KBM (Algorithm 1) within SMC, we need to consider the times at which to conduct resampling. One could in principle devise a scheme to resample in continuous time, however for simplicity and theoretical reasons we only conduct resampling at pre-specified times $t_0 := 0 < t_1 < \ldots < t_m := T$ (we would naturally choose these such that $t_i - t_{i-1} := T/m$). This is presented as the *Quasi-Stationary Monte Carlo (QSMC)* algorithm in Algorithm 2 in which we define N_{eff} as the effective sample size of [34], and N_{th} as a user chosen threshold for which will conduct (on a given iteration) resampling of the particle set at the resampling time. The algorithm outputs the weighted particles at the end of each iteration.

Given the output from Algorithm 2, we can then estimate the target distribution π. We first choose a burn-in time, t^*, such that we believe we have achieved convergence to the quasi-stationary distribution of killed Brownian motion. Our

Algorithm 2 Quasi-Stationary Monte Carlo Algorithm (QSMC) Algorithm.

1. **Initialisation Step** ($i = 0$)
   (a) Input: Starting value, $\hat{\mathbf{x}}$, number of particles, N, and set of m times $t_{1:m}$.
   (b) $\mathbf{X}_0^{(\cdot)}$: For k in 1 to N set $\mathbf{X}_{t_0}^{(1:N)} = \hat{\mathbf{x}}$ and $w_{t_0}^{(1:N)} = 1/N$.
2. **Iterative Update Steps** ($i = i + 1$ **while** $i \leq m$)
   (a) N_{eff}: If $N_{\text{eff}} \leq N_{\text{th}}$ then for k in 1 to N resample $\mathbf{X}_{t_{i-1}}^{(k)} \sim \tilde{\pi}_{t_{i-1}}^N$, the empirical distribution defined by the current set of weighted particles, and set $w_{t_{i-1}}^{(k)} = 1/N$.
   (b) For k in 1 to N,
      i. $\mathbf{X}_{t_i}^{(\cdot)}$ Simulate $\mathbf{X}_{t_i}^{(k)} | \mathbf{X}_{t_{i-1}}^{(k)}$ along with un-normalised weight increment $w_{t_i - t_{i-1}}^*$ as per Algorithm 1.
      ii. $w_{t_i}^{\prime(\cdot)}$: Calculate un-normalised weights, $w_{t_i}^{\prime(k)} = w_{t_{i-1}}^{(k)} \cdot w_{t_i - t_{i-1}}^*$.
   (c) $w_{t_i}^{(\cdot)}$: For k in 1 to N set $w_{t_i}^{(k)} = w_{t_i}^{\prime(k)} / \sum_{l=1}^N w_{t_i}^{\prime(l)}$.
   (d) $\tilde{\pi}_{t_i}^N$: Set $\tilde{\pi}_{t_i}^N(\mathrm{d}\mathbf{x}) := \sum_{k=1}^N w_{t_i}^{(k)} \cdot \delta_{\mathbf{X}_{t_i}^{(k)}}(\mathrm{d}\mathbf{x})$.

approximation to the law of the killed process is then simply the weighted occupation measures of the particle trajectories in the interval $[t^*, T]$. More precisely, using the output of the QSMC algorithm then

$$\pi(\mathrm{d}\mathbf{x}) \approx \hat{\pi}(\mathrm{d}\mathbf{x}) := \frac{1}{|\{i : t_i \in [t^*, T]\}|} \sum_{i=1 : t_i \in [t^*, T]}^{m} \sum_{k=1}^{N} w_{t_i}^{(k)} \cdot \delta_{\mathbf{X}_{t_i}^{(k)}}(\mathrm{d}\mathbf{x}). \tag{29}$$

In Appendix H we use existing theory for SMC to obtain a central limit theorem for estimators based on the output QSMC. We defer the theoretical characterisation of the SMC approach used to Section 4.2, at which point we can study QSMC with full application to our big data context. Provided the timescale upon which resampling is conducted is not too coarse, and low variance resampling strategies are employed (for instance, [20]), one would not expect qualitative differences in performance between the approaches.

4. Sub-sampling

In Section 3 we presented two methodological implementations for simulating our desired quasi-limiting distribution – in particular, we simulated a number of trajectories of Brownian motion which were either killed or re-weighted after an exponentially distributed amount of time (with state space dependent intensity). As we will outline, in the case of a big data setting the most naïve implementations of the killing and re-weighting mechanisms are the $\mathcal{O}(n)$ bottlenecks that we need to overcome.

We begin here by studying a simple property of our estimator of $P(\mathbf{X})$, the

survival probability of a Brownian motion path. In Section 4.1 we exploit this property to overcome the $\mathcal{O}(n)$ discussed above.

To recap, in Section 3.1 both our rejection sampler (KBM) and our importance sampler (IS-KBM) used the estimator

$$P(\mathbf{X}) = \prod_{i=1}^{n_R} \Bigg[\underbrace{\exp\left\{(\Phi - L_{\mathbf{X}}^{(i)}) \cdot [(\tau_i \wedge T) - \tau_{i-1}]\right\}}_{=:P^{(1)}(\mathbf{X}) \in [0,1]} \cdot \mathbb{E}_{\mathbb{U}_R^{(i)}} \Bigg[\prod_{j=1}^{\kappa_i} \underbrace{\left(\frac{U_{\mathbf{X}}^{(i)} - \phi\left(\mathbf{X}_{\xi_j^{(i)}}\right)}{U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)}} \right)}_{=:P^{(2,i,j)}(\mathbf{X}) \in [0,1]} \Bigg| \mathbf{X} \Bigg] \Bigg], \tag{30}$$

where the outer product is over the different layers, the inner product is over the events within a layer, and the expectation is with respect to the number and position of these events (for which we have observed the proposal Brownian motion sample path). Note that the $\mathcal{O}(n)$ bottleneck within this estimator is the evaluation of $\phi(\cdot)$.

We can replace each term within the inner product, $P^{(2,i,j)}(\mathbf{X})$, by an unbiased estimator, and this identity will still hold in expectation. That is we introduce an auxiliary random variable $A \sim \mathcal{A}$, and construct estimator $\tilde{\phi}_A(\cdot)$ such that

$$\mathbb{E}_{\mathcal{A}}\left[\tilde{\phi}_A(\cdot)\right] = \phi(\cdot). \tag{31}$$

Recall that to ensure in (30) that $P(\mathbf{X}) \in [0,1]$, we used the bounds provided by each layer to find $L_{\mathbf{X}}^{(i)}, U_{\mathbf{X}}^{(i)} \in \mathbb{R}$ such that $\phi(\mathbf{X}_{[\tau_{i-1},\tau_i]}) \in [L_{\mathbf{X}}^{(i)}, U_{\mathbf{X}}^{(i)}]$. To ensure $P(\mathbf{X}) \in [0,1]$ on replacement of the terms in the inner product, $P^{(2,i,j)}(\mathbf{X})$, we would require new bounds $\tilde{\Phi} \leq \tilde{L}_{\mathbf{X}}^{(i)} \leq L_{\mathbf{X}}^{(i)}$ and $\tilde{U}_{\mathbf{X}}^{(i)} \geq U_{\mathbf{X}}^{(i)}$ such that $\forall A \sim \mathcal{A}$ we have $\tilde{\phi}_A(\mathbf{X}_{[\tau_{i-1},\tau_i]}) \in [\tilde{L}_{\mathbf{X}}^{(i)}, \tilde{U}_{\mathbf{X}}^{(i)}]$. Note under this modification we would have the same number of layers, n_R, but the number of events in each layer would increase. The resulting modified estimator, $P(\mathbf{X})$, is as follows

$$P(\mathbf{X}) = \prod_{i=1}^{n_R} \Bigg[\underbrace{\exp\left\{(\tilde{\Phi} - \tilde{L}_{\mathbf{X}}^{(i)}) \cdot [(\tau_i \wedge T) - \tau_{i-1}]\right\}}_{=:P^{(1)}(\mathbf{X}) \in [0,1]} \cdot \mathbb{E}_{\tilde{\mathbb{U}}_R^{(i)}, \mathcal{A}} \Bigg[\prod_{j=1}^{\tilde{\kappa}_i} \underbrace{\left(\frac{\tilde{U}_{\mathbf{X}}^{(i)} - \tilde{\phi}_{A_j}\left(\mathbf{X}_{\xi_j^{(i)}}\right)}{\tilde{U}_{\mathbf{X}}^{(i)} - \tilde{L}_{\mathbf{X}}^{(i)}} \right)}_{=:P^{(2,i,j)}(\mathbf{X}) \in [0,1]} \Bigg| \mathbf{X} \Bigg] \Bigg], \tag{32}$$

The expectation in (32) is now also with respect to the auxiliary random variables, and we need to use independent copies of A for each time we calculate the

estimator $\tilde{\phi}_A$. Proving this identity is straight-forward, as we just take expectations with respect to the auxiliary random variables, and use the unbiasedness property of $\tilde{\phi}_A$, to get back to (30) (albeit with a different choice of lower and upper bound). One should note that (30) holds for all valid lower and upper bounds, including for those chosen in (32). As such, one could study the direct effect of the unbiased estimator replacement in (31) by simply contrasting (30) and (32) using the same bounds $(\tilde{L}_{\mathbf{X}}^{(i)}, \tilde{U}_{\mathbf{X}}^{(i)})$.

4.1. Constructing a scalable replacement estimator

Consider the estimator we developed in (32). If we construct our auxiliary random variable $A \sim \mathcal{A}$ and estimator $\tilde{\phi}_A(\cdot)$ with $\mathcal{O}(1)$ computational complexity, then, if for a given layer we define $\tilde{\lambda} := \tilde{U}_{\mathbf{X}}^{(i)} - \tilde{L}_{\mathbf{X}}^{(i)}$ (which will determine the intensity of events in the i^{th} layer), the computational cost of the modified algorithm is simply proportional to $\tilde{\lambda}$. Note that the intensity of events will increase relative to the estimator in (30) as $\lambda := U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)} \leq \tilde{\lambda}$. Therefore there is a trade-off between the computational loss in this regard, with that saved from employing the unbiased estimator $\tilde{\phi}_A(\cdot)$.

Recalling the construction of KBM in Section 3,

$$\phi(\mathbf{x}) := (||\mathbf{\Lambda}^{1/2}\nabla \log \pi(\mathbf{x})||^2 + \operatorname{div} \mathbf{\Lambda}^{1/2}\nabla \log \pi(\mathbf{x}))/2 - \Phi. \tag{33}$$

Developing an unbiased estimator with $\mathcal{O}(1)$ complexity to determine acceptance is trivial. For instance considering a single component of (33), $\mathbf{\Lambda}^{1/2}\nabla \log \pi(\mathbf{x})$, then (as outlined in Section 1 and as motivated by [60]) the naïve estimator with complexity $\mathcal{O}(1)$ is simply $(n+1)\mathbf{\Lambda}^{1/2}\nabla \log f_I(\mathbf{x})$, where $I \sim \mathrm{U}\{0, \ldots, n\}$. Unfortunately our naïve estimator of $\mathbf{\Lambda}^{1/2}\nabla \log \pi(\mathbf{x})$ has a bound which grows at least as $\mathcal{O}(n^{1/2})$. Therefore the computational speed-up of employing an $\mathcal{O}(1)$ unbiased subsampling estimator of (33) would be counteracted by having a $\tilde{\lambda}$ (and consequently an increase in the intensity of the events required in (32)) which grows as $\mathcal{O}(n)$. By comparison, the bound achieved when not employing the sub-sampling estimator, λ, will be $\mathcal{O}(n)$ times smaller than $\tilde{\lambda}$.

The argument we develop in the remainder of this section is that by employing a control variate approach for the construction of the unbiased estimator, $\tilde{\phi}_A(\cdot)$, we can control the computational growth of the corresponding $\tilde{\lambda}$ with n.

We begin by choosing a point close to a mode of the posterior distribution π, denoted by $\hat{\mathbf{x}}$. In fact for our scaling arguments to hold, we require $\hat{\mathbf{x}}$ to be within $\mathcal{O}(n^{-1/2})$ of the true model, and achieving this is a less demanding task than actually locating the mode. Moreover we note that this operation is only required to be done once, and not at each iteration.

Letting $\mathcal{A}$ be the law of $I \sim \mathrm{U}\{0,\ldots,n\}$ we have

$$\mathbb{E}_{\mathcal{A}}\big[\underbrace{(n+1)\cdot[\nabla \log f_I(\mathbf{x}) - \nabla \log f_I(\hat{\mathbf{x}})]}_{=:\tilde{\alpha}_I(\mathbf{x})}\big] = \underbrace{\nabla \log \pi(\mathbf{x}) - \nabla \log \pi(\hat{\mathbf{x}})}_{=:\alpha(\mathbf{x})}. \quad (34)$$

As such, we can re-express $\phi(\mathbf{x})$ as

$$\phi(\mathbf{x}) = (\alpha(\mathbf{x})^T \mathbf{\Lambda}\,(2\nabla \log \pi(\hat{\mathbf{x}}) + \alpha(\mathbf{x})) + \mathbf{\Lambda}^{1/2}\mathrm{div}\,\alpha(\mathbf{x}))/2 + C, \quad (35)$$

where $C := (||\mathbf{\Lambda}^{1/2}\nabla \log \pi(\hat{\mathbf{x}})||^2 + \mathrm{div}\,\mathbf{\Lambda}^{1/2}\nabla \log \pi(\hat{\mathbf{x}}))/2 - \Phi$ is a constant. Letting $\mathcal{A}$ now be the law of $I, J \overset{\text{iid}}{\sim} \mathrm{U}\{0,\ldots,n\}$ we can construct the following unbiased estimator of ϕ,

$$\mathbb{E}_{\mathcal{A}}\big[\underbrace{(\tilde{\alpha}_I(\mathbf{x}))^T \mathbf{\Lambda}\,(2\nabla \log \pi(\hat{\mathbf{x}}) + \tilde{\alpha}_J(\mathbf{x})) + \mathbf{\Lambda}^{1/2}\mathrm{div}\,\tilde{\alpha}_I(\mathbf{x}))/2 + C}_{=:\tilde{\phi}_A(\mathbf{x})}\big] = \phi(\mathbf{x}). \quad (36)$$

The construction of our estimator requires evaluation of the constants $\nabla \log \pi(\hat{\mathbf{x}})$ and $\mathrm{div}\,\nabla \log \pi(\hat{\mathbf{x}})$. Although both are be $\mathcal{O}(n)$ evaluations they only have to be computed once, and furthermore their calculation can be done entirely in parallel. For a given target density we would then proceed to compute a suitable $\tilde{\lambda}$ for our sub-sampler. We return at the end of this section to consider finding an appropriate $\tilde{\lambda}$ and its growth with n.

Embedding our sub-sampling estimator described above within the QSMC algorithm of Section 3.2, we arrive at Algorithm 3 which we call the *Scalable Langevin Exact algorithm (ScaLE)*. A similar modification could be made to the rejection-sampling version, R-QSMC, which we discussed in Section 3.2 and detailed in Appendix F. We term this variant *Rejection Scalable Langevin Exact algorithm (R-ScaLE)* and provide full algorithmic details in Appendix G.

Algorithm 3 The ScaLE Algorithm (as per Algorithm 2 unless stated otherwise).

0. Choose $\hat{\mathbf{x}}$ and compute $\nabla \log \pi(\hat{\mathbf{x}})$, $\mathrm{div}\,\nabla \log \pi(\hat{\mathbf{x}})$.

2(b)i. On calling Algorithm 1,

(a) Replace $L^{(i)}_{\mathbf{X}}, U^{(i)}_{\mathbf{X}}$ in Step 2 with $\tilde{L}^{(i)}_{\mathbf{X}}, \tilde{U}^{(i)}_{\mathbf{X}}$.

(b) Replace Step 7 with: τ_i: If $\xi_j = \tau_i$, set $i = i+1$, and return to Step 2. Else simulate $A_j = (I_j, J_j)$, with $I_j, J_j \overset{\text{iid}}{\sim} \mathrm{U}\{0,\ldots,n\}$, and set $w^*_{\xi_j} = w^*_{\xi_j} \cdot (\tilde{U}^{(i)}_{\mathbf{X}} - \tilde{\phi}_{A_j}(\mathbf{X}_{\xi_j}))/(\tilde{U}^{(i)}_{\mathbf{X}} - \tilde{L}^{(i)}_{\mathbf{X}})$ (where $\tilde{\phi}_{A_j}$ is defined as in (36)) and return to Algorithm 1 Step 3.

As previously motivated, given our auxiliary random variable $A \sim \mathcal{A}$ and estimator $\tilde{\phi}_A(\cdot)$ both have $\mathcal{O}(1)$ computational complexity, the iterative computational complexity and scalability of the ScaLE algorithm is determined entirely by $\tilde{\lambda}$.

In particular, we will analyse the computational cost of ScaLE by studying the growth of the bound on (36) as a function of n.

For simplicity, the argument we present assumes that we have a bound on the contribution to the second-derivative of the log-likelihood for each datum. That is we assume $\forall I \in \{0, \dots, n\}$

$$|\nabla \nabla^T \log f_I(\mathbf{x})| \leq K_n, \tag{37}$$

for some $K_n > 0$. Frequently this bound will hold globally. For example, considering densities with Gaussian and heavier tails, then the Hessian of the log-likelihood is typically uniformly bounded (in both data and parameter). However, as the argument we present below illustrates, we actually only require (37) to hold locally within a bounded region. We return at the end of this section to discuss the dependence of K_n on the size of the data set, n.

Recall from Section 3.1 that the practitioner is able to select an appropriate diagonal pre-conditioning matrix, which we write in the form $\mathbf{\Lambda} = n^{-1}\mathbf{\Lambda}_0$, where $\mathbf{\Lambda}_0$ approximately homogenises the scale of the components. We assume such a scaling of the preconditioning matrix as the size of the movements of our killed Brownian motion over a fixed time-interval will then be of $O(n^{-1/2})$, and hence contract at the same rate as the posterior, at least in the regular inference case.

Recalling the layer construction of Section 3.1 for a single trajectory of killed Brownian motion, we can ensure that over any finite time interval we have $\mathbf{x} \in R_{\mathbf{X}} := \mathbf{\Lambda}^{1/2}[\ell_{\mathbf{X}}, \upsilon_{\mathbf{X}}]$ (where $[\ell_{\mathbf{X}}, \upsilon_{\mathbf{X}}]$ denotes the hypercuboid with given corners). We shall set $\mathbf{x}^* = \mathbf{\Lambda}^{1/2}(\ell_{\mathbf{X}} + \upsilon_{\mathbf{X}})/2$.

The practitioner has complete freedom to choose $R_{\mathbf{X}}$, and following the guidance given in (67) of Appendix C.3), it makes sense that $\ell_{\mathbf{X}}$ and $\upsilon_{\mathbf{X}}$ be $\mathcal{O}(1)$, which in turn ensures that $||\mathbf{x} - \mathbf{x}^*|| < Cn^{-1/2}$ for some $C > 0$ and for all $\mathbf{x} \in R_{\mathbf{X}}$.

As we are interested in the magnitude of $\tilde{\lambda}$ for $\mathbf{x} \in R_{\mathbf{X}}$, this can be set to be twice the maximum of $|\tilde{\phi}_A(\mathbf{x})|$ for $\mathbf{x} \in R_{\mathbf{X}}$, over all possible realisations of A. To bound $|\tilde{\phi}_A(\mathbf{x})|$, we begin by first considering our elementary estimator in (34). By imposing our condition in (37) we have

$$\max_{\mathbf{x} \in R_{\mathbf{X}}, I \in \{0,\dots,n\}} |\tilde{\alpha}_I(\mathbf{x})| \leq (n+1) \cdot K_n \cdot \max_{\mathbf{x} \in R_{\mathbf{X}}} ||\mathbf{x} - \hat{\mathbf{x}}||. \tag{38}$$

We can proceed to find a bound for our estimator in (36) as follows

$$\begin{aligned}\max_{\mathbf{x}\in R_{\mathbf{X}},A\in\mathcal{A}}|\tilde{\phi}_A(\mathbf{x})| \leq K_n\cdot n^{-1}\cdot(n+1)\cdot\max_{\mathbf{x}\in R_{\mathbf{X}}}||\mathbf{x}-\hat{\mathbf{x}}||\cdot||\mathbf{\Lambda}_0|| \\ \cdot(2\nabla\log\pi(\hat{\mathbf{x}})+(n+1)\cdot K_n\cdot\max_{\mathbf{x}\in R_{\mathbf{X}}}||\mathbf{x}-\hat{\mathbf{x}}||)/2 \\ + K_n\cdot n^{-1/2}\cdot||\mathbf{\Lambda}_0^{1/2}||\cdot\mathrm{div}\left((n+1)\cdot\max_{\mathbf{x}\in R_{\mathbf{X}}}||\mathbf{x}-\hat{\mathbf{x}}||\right)/2+C.\end{aligned} \tag{39}$$

Now, as we have $\max_{\mathbf{x}\in R_{\mathbf{X}}}||\mathbf{x}-\hat{\mathbf{x}}|| \leq ||\mathbf{x}^*-\hat{\mathbf{x}}||+Cn^{-1/2}$, then

$$\begin{aligned}\tilde{\lambda} = \mathcal{O}\Big(K_n\cdot\left(||\mathbf{x}^*-\hat{\mathbf{x}}||+Cn^{-1/2}\right)\cdot\left(1+nK_n\cdot\left(||\mathbf{x}^*-\hat{\mathbf{x}}||+Cn^{-1/2}\right)\right) \\ + n^{1/2}\cdot K_n\cdot\left(||\mathbf{x}^*-\hat{\mathbf{x}}||+Cn^{-1/2}\right)\Big).\end{aligned} \tag{40}$$

At stationarity $\mathbf{x}^*$ will be a draw from the support of the posterior. So, if we assume posterior contraction at the regular $\sqrt{n}$ rate, then $||\mathbf{x}^*-\hat{\mathbf{x}}|| = \mathcal{O}_p(n^{-1/2})$, and hence $\tilde{\lambda} = \mathcal{O}_p(1)$.

Now, returning to consider typical values of K_n, we first note that for smooth densities with Gaussian and heavier tails, the Hessian of the log-likelihood is typically uniformly bounded (in both data and parameter). As such it is automatic that K_n should be constant in this case and so the ScaLE algorithm has $\mathcal{O}(1)$ iterative computational complexity.

For light-tailed likelihoods, the size of K_n will likely depend upon the most "extreme" data. Since maxima of heavy-tailed distributions are typically $\mathcal{O}((\log n)^\eta)$ for some constant $\eta \leq 1/2$, this can be expected to affect K_n making it grow logarithmically as a function of n. We shall omit detailed calculations here. This discussion does however suggest that non-linear transformations could potentially be used to completely eliminate this problem should it be considered important.

Example: As a toy example to illustrate the effect of subsampling and control variates, we carried out an analysis of the Bayesian posterior for a t_5 model

$$\pi(\theta|\mathbf{y}) \propto \prod_{i=1}^{n}\left(\frac{1}{5+(y_i-\theta)^2}\right)^3, \tag{41}$$

with data generated from the same model with $\theta = 0$. Figure 2 empirically illustrates the scaling properties of the approach, showing that the for this example the computational cost per iteration remains more or less constant as a function of n, as predicted by the above arguments, so that there is essentially no iterative cost associated with n.

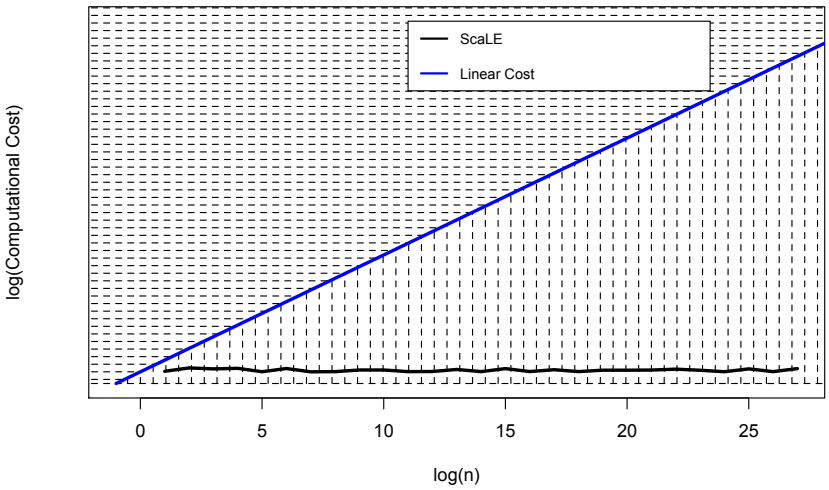

FIGURE 2. *Iterative computational cost of ScaLE (black line) when applied to increasing volumes of data generated from a one dimensional t_5 distribution, with algorithmic parameters chosen to give a similar overall Monte Carlo error for each data-set size. Blue line depicts linear computational cost, vertical hatching sub-linear computational cost, and horizontal hatching super-linear computational cost.*

4.2. Theoretical Properties

SMC algorithms in both discrete and continuous time have been studied extensively in the literature. In order to avoid a lengthy technical diversion, we restrict ourselves here to studying a slightly simplified version of the problem in order to obtain the simplest and most interpretable possible form of results. We defer the technical details of this construction to Appendix H and give here only a qualitative description intended to guide intuition and the key result: that the resulting estimator satisfies a Gaussian central limit theorem with the usual Monte Carlo rate.

We consider a variant of the algorithm in which (multinomial) resampling occurs at times kh for $k \in \mathbb{N}$ where h is a timestep specified in advance and consider the behaviour of estimates obtained at these times. Extension to resampling at a random subset of these resampling times would be possible using the approach of [42], considering precisely the QSMC algorithm presented in Algorithm 2 and the ScaLE algorithm in Algorithm 3 would require additional technical work somewhat beyond the scope of this paper; we did not observe any substantial difference in behaviour.

In order to employ standard results for SMC algorithms it is convenient to consider a discrete time embedding of the algorithms described. We consider an abstract formalism in which between the specified resampling times the trajectory of the Brownian motion is sampled, together with such auxiliary random variables as are required in any particular variant of the algorithm. Provided the potential function employed to weight each particle prior to resampling has conditional expectation (given the path) proportional to the exact killing rate integrated over these discrete time intervals we recover a valid version of the ScaLE algorithm.

This discrete time formalism allows for results on more standard SMC algorithms to be applied directly to the ScaLE Algorithm. We provide in the following proposition a straightforward corollary to [41, Chapter 9], which demonstrates that estimates obtained from a single algorithmic time slice of the ScaLE algorithm satisfy a central limit theorem.

Proposition 1 (Central Limit Theorem)**.** *In the context described, under mild regularity conditions (see references given in Appendix H):*

$$\lim_{N\to\infty} \sqrt{N}\left[\frac{1}{N}\sum_{i=1}^{N}\varphi(X_{hk}^i) - \mathbb{E}_{\mathbb{K}_{hk}^x}\left[\varphi(X_{hk}^i)\right]\right] \Rightarrow \sigma_k(\varphi)Z$$

where, $\varphi : \mathbb{R}^d \to \mathbb{R}$, Z is a standard normal random variable, $\Rightarrow$ denotes convergence in distribution, and $\sigma_k(\varphi)$ depends upon the precise choice of subsampling scheme as well as the test function of interest and is specified in Appendix H.

5. Examples

In this section we present two applications of the methodology developed in this paper. In particular, in Section 5.1 for exposition we consider the application of ScaLE to a simple example of logistic regression with a moderately large data set. In Section 5.2 we consider a problem in which we wish do parameter inference for a more complicated mixture model, motivated by big data application and with $n = 2^{27}$.

5.1. Example 1: Menarche

In this subsection we consider the R [56] dataset 'Menarche', due originally to [39], which appears in the 'MASS' package [57]. Menarche is a dataset comprising the age at which $n = 3918$ adolescent female children reach menarche. This dataset is frequently used as a simple textbook example to which we can fit a single feature logistic regression model. We will contrast the likelihood-based estimates provided by the R glm function for this dataset, with those provided by

the ScaLE algorithm (Algorithm 3) we developed in Section 4.1. More precisely, we have binary observations of and fit a model of the form,

$$y_i = \begin{cases} 1 & \text{with probability } \dfrac{\exp\{\mathbf{x}_i^T\beta\}}{1+\exp\{\mathbf{x}_i^T\beta\}}, \\ 0 & \text{otherwise,} \end{cases} \tag{42}$$

noting that we can readily compute for each datum its contribution to the likelihood along with its derivatives. We normalise the design matrix (subtracting the sample mean and dividing by standard deviation); the glm [R] package estimates $\hat{\beta} = [1.410, 4.658]^T$. We show the glm estimate as the blue line in Figure 3.

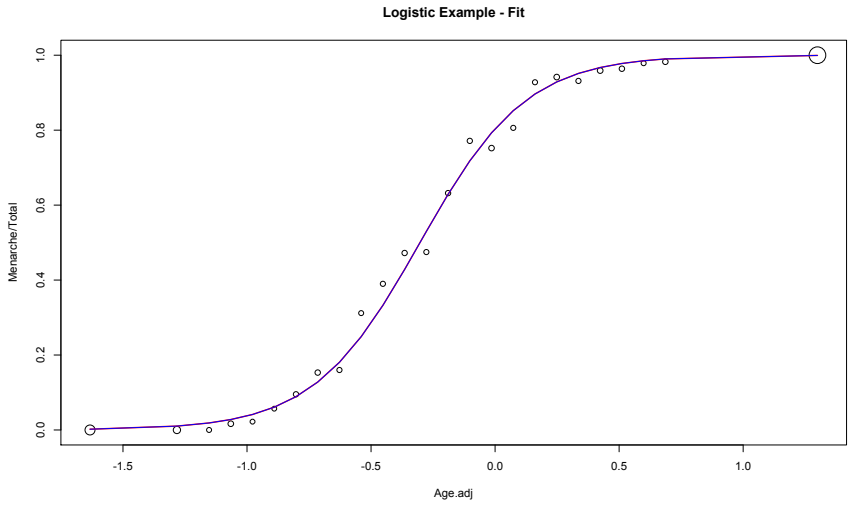

FIGURE 3. *Contrast of glm (blue) and ScaLE (red) fit to the MASS Menarche dataset.*

In contrast, we initialise the ScaLE algorithm at a point one standard deviation from $\hat{\beta}$, with a particle set of size $N = 100$. In total we run ScaLE for diffusion time of $T = 1000$, and approximate the law of the killed process as in (29) using observations of the trajectory at a resolution of $t_i - t_{i-1} = 0.1$. Each particle trajectory at each time $t \in [0, T]$ is associated with a sub-sample of the full dataset of size 2. The contrast with glm of the fit to the dataset is shown as the red line in Figure 3. An example of a typical run can be found in Figure 4. In Figure 5 we compute the profile of the log-likelihood along with the estimates of glm and ScaLE, contrasting the profile with a kernel density estimate computed using the locations of the weighted trajectories at the terminal time slice (as in (28)).

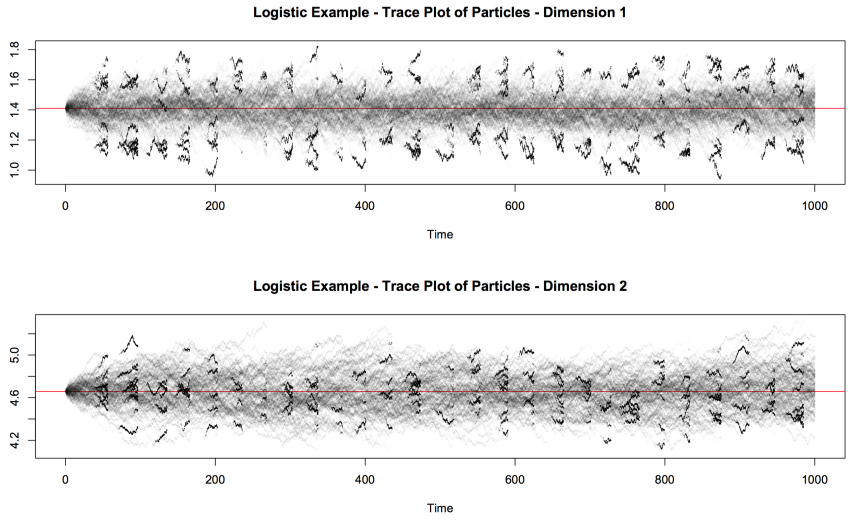

FIGURE 4. *Particle trajectories of ScaLE applied to the Menarche dataset.*

5.2. Example 2: Contaminated Mixture

In this subsection we consider the application of ScaLE to conduct parameter inference for a contaminated mixture model. This is motivated by big data sets obtained from internet applications in which the dataset obtained is corrupted with noisy observations. In particular, we the case in which each datum comprises two features and we consider fitting a model in which the likelihood of an individual observation (y_i) is as follows,

$$F_i := \frac{1-p}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{1}{2\sigma^2}\left(\alpha \cdot x_{i,1} + \beta \cdot x_{i,2} - y_i\right)^2\right\} + \frac{p}{\sqrt{2\pi\phi^2}} \exp\left\{-\frac{1}{2\phi^2} y_i^2\right\}. \tag{43}$$

In this model p represents the level of corruption and ϕ the variance of the corruption. A common approach if we wished to conduct parameter inference for this model using MCMC is data augmentation [54]. However, for large data sets this is not feasible as the dimensional of the auxiliary variable vector will be $\mathcal{O}(n)$.

We generate a dataset of size $n = 2^{27}$ from the model with parameters $\mu = [\alpha, \beta, \sigma, \phi, p] = [2, 5, 1, 10, 0.05]$, and conduct parameter inference using ScaLE initialised at a point several deviations from that used to generate the data. For this application we use a particle set of size $N = 2048$, run ScaLE for diffusion time of $T = 200$, and approximate the law of the killed process as in (29) using observations of the trajectory at a resolution of $t_i - t_{i-1} = 0.1$. Each particle trajectory at each time $t \in [0, T]$ is associated with a sub-sample of the full

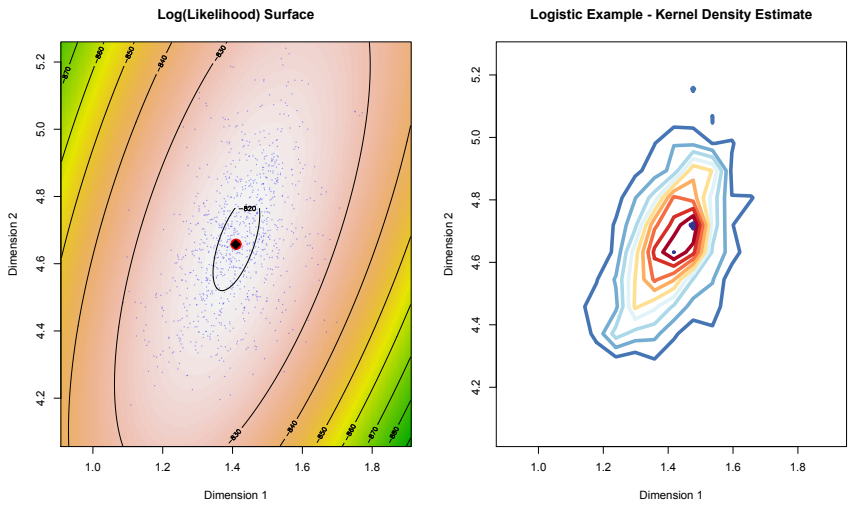

FIGURE 5. *Log-likelihood surface and ScaLE estimate for the Menarche dataset.*

dataset of size 32. An example of a typical run can be found in Figure 6, and in Figure 7 estimate the marginal density of the parameters using the occupation measure of the trajectories in the interval $t \in [100, 200]$.

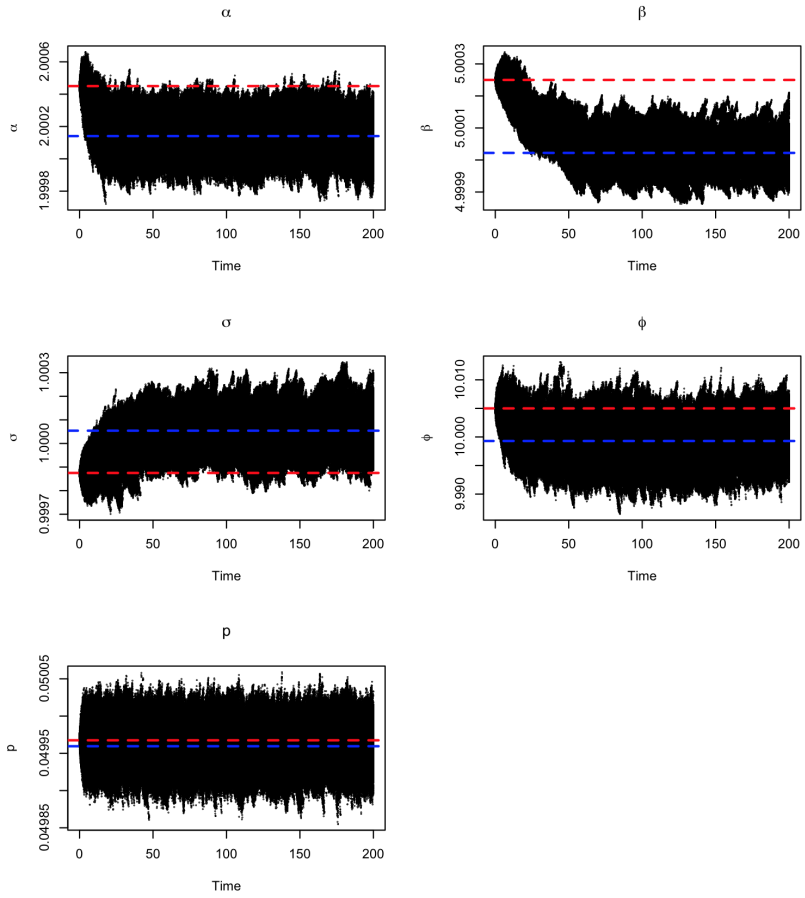

FIGURE 6. *Particle trajectories of ScaLE applied to the contaminated mixture dataset. Red lines mark the parameter values used to initialise the algorithm, blue lines mark the mode of the marginal density.*

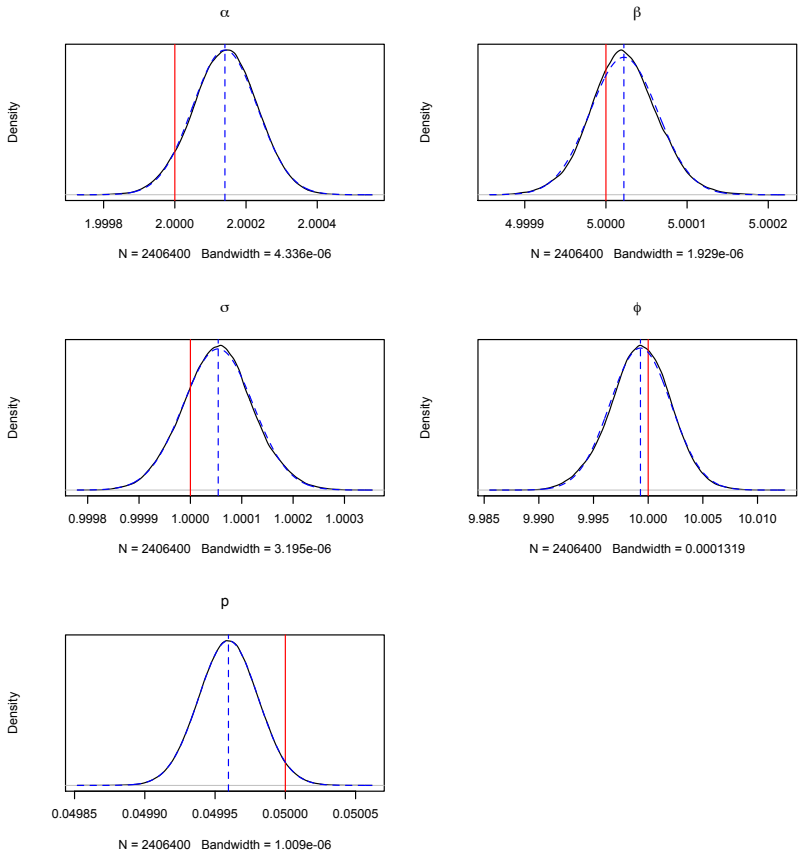

FIGURE 7. *Marginal densities obtained by ScaLE for the contaminated mixture dataset. Red lines mark the parameter values for which the data set was generated.*

6. Conclusions

In this paper we have introduced a new class of *Quasi-Stationary Monte Carlo (QSMC)* methods which are genuinely continuous-time algorithms for simulating from complex target distributions. We have emphasised its particular effectiveness in the context of *big data* by developing novel sub-sampling approaches and the *Scalable Langevin Exact (ScaLE)* algorithm. Unlike its immediate competitors, our subsampling approach within ScaLE is essentially computationally free and does not result in any approximation to the target distribution. Our methodology is embedded within an SMC framework, which we use to provide theoretical results. In addition we give examples to which we apply ScaLE, demonstrating its robust scaling properties for large data sets.

We see this paper as a first step in what we hope will be a fruitful new direction for Computational Statistics. Many ideas for variations and extensions our implementation exist and will stimulate further investigation.

Firstly, the need to simulate a quasi-stationary distribution creates particular challenges. Although quasi-stationarity is underpinned by an elegant mathematical theory, the development of numerical methods for quasi-stationarity is understudied. We have presented an SMC methodology for this problem, but alternatives exist. For instance, [10] suggest alternative approaches.

Even within an SMC framework for extracting the quasi-stationary distribution, there are interesting alternatives we have not explored. For example, by a modification of our re-weighting mechanism we can relate the target distribution of interest to the limiting *smoothing* distribution of the process, as opposed to the filtering distribution as we do here. Within the quasi-stationary literature this is often termed the type II quasi-stationary distribution. As such, the rich SMC literature offers many other variations on the procedures we adopt here.

We have also concentrated solely on killed (or re-weighted) Brownian motion in our paper. We have shown that this strategy has robust convergence properties. However, given existing methodology for the exact simulation of diffusions in [8, 7, 48], there is scope to develop methods which use proposal measures which much better *mimick* the shape of the posterior distribution.

Our sub-sampling and control variate approaches offer dramatic computational savings. However there may be scope to extend these ideas still further. For instance, it might be possible to sub-sample dimensions and thus reduce the dimensional complexity for implementing each iteration.

We conclude by noting that as a byproduct, the theory behind our methodology offers new insights into problems concerning the existence of quasi-stationary distributions for diffusions killed according to a state-dependent hazard rate, complementing and extending current state-of-the-art literature [53].

Acknowledgments

MP was supported by the EPSRC [grant number EP/K014463/1]. PF was supported by [grant numbers EP/N031938/1, EP/K014463/1]. GOR was supported by [grant numbers EP/K034154/1, EP/K014463/1, EP/D002060/1].

References

[1] Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Annals of Statistics* **37** 697–725.

[2] Asmussen, S., Glynn, P. and Pitman, J. (1995). Discretization error in simulation of one-dimensional reflecting Brownian motion. *Annals of Applied Probability* **5** 875–896.

[3] Bardenet, R., Doucet, A. and Holmes, C. (2014). Towards scaling up Markov chain Monte Carlo: an adaptive subsampling approach. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)* (T. Jebara and E. P. Xing, eds.) 405-413. JMLR Workshop and Conference Proceedings.

[4] Bardenet, R., Doucet, A. and Holmes, C. (2015). On Markov chain Monte Carlo methods for tall data. *arXiv preprint arXiv:1505.02827*.

[5] Besag, J. (1994). Comments on "Representations of knowledge in complex systems" by U. Grenander and M. Miller. *Journal of the Royal Statistical Society B* **56** 591–592.

[6] Beskos, A., Papaspiliopoulos, O. and Roberts, G. O. (2006). Retrospective Exact Simulation of Diffusion Sample Paths with Applications. *Bernoulli* **12** 1077–1098.

[7] Beskos, A., Papaspiliopoulos, O. and Roberts, G. O. (2008). A Factorisation of Diffusion Measure and Finite Sample Path Constructions. *Methodology and Computing in Applied Probability* **10** 85–104.

[8] Beskos, A. and Roberts, G. O. (2005). An Exact Simulation of Diffusions. *Annals of Applied Probability* **15** 2422–2444.

[9] Bierkens, J., Fearnhead, P. and Roberts, G. O. (2016). The Zig-Zag Process and Super-Efficient Sampling for Bayesian Analysis of Big Data. *arXiv preprint arXiv:1607.03188*.

[10] Blanchet, J. H., Glynn, P. and Zheng, S. (2015). Theoretical analysis of a Stochastic Approximation approach for computing Quasi-Stationary distributions of general state space Markov chains. *arXiv preprint arXiv:1502.06451*.

[11] Burq, Z. A. and Jones, O. (2008). Simulation of Brownian motion at first passage times. *Mathematics and Computers in Simulation* **77** 64–71.

[12] Chen, N. and Huang, Z. (2013). Localisation and Exact Simulation of Brownian Motion Driven Stochastic Differential Equations. *Mathematics of Operational Research* **38** 591–616.

[13] Chen, T., Fox, E. B. and Guestrin, C. (2014). Stochastic Gradient Hamiltonian Monte Carlo. In *ICML* 1683–1691.

[14] Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its applications to Bayesian inference. *Annals of Statistics* **32** 2385–2411.
[15] Ciesielski, Z. and Taylor, S. J. (1962). First passage times and sojourn times for Brownian motion in space and the exact Hausdorff measure of the sample path. *Annals of Mathematical Statistics* **103** 1434–450.
[16] Collet, P., Martinez, S. and San Martin, S. (2013). *Quasi-Stationary Distributions: Markov Chains, Diffusions and Dynamical Systems.* Springer, Berlin.
[17] Dachuna-Castelle, D. and Florens-Zmirou, D. (1986). Estimation of the coefficients of a diffusion from discrete observations. *Stochastics* **19** 263–284.
[18] Devroye, L. (1986). *Non-Uniform Random Variate Generation.*, 1st ed. Springer, New York.
[19] Devroye, L. (2009). On exact simulation algorithms for some distributions related to Jacobi theta functions. *Statistics and Probability Letters* **79** 2251–2259.
[20] Douc, R., Cappé, O. and Moulines, E. (2005). Comparison of resampling schemes for particle filters. In *Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis* **I** 64-69. IEEE.
[21] Doucet, A. and Johansen, A. M. (2011). A Tutorial on Particle Filtering and Smoothing: Fifteen years later. In *The Oxford Handbook of Nonlinear Filtering* (D. Crisan and B. Rozovsky, eds.) 656–704. Oxford University Press.
[22] Durmus, A. and Moulines, E. (2015). Non-asymptotic convergence analysis for the Unadjusted Langevin Algorithm. *ArXiv e-prints 1507.05021.*
[23] Fort, G. and Roberts, G. O. (2005). Subgeometric ergodicity of strong Markov processes. *The Annals of Applied Probability* **15** 1565–1589.
[24] Gordon, N. J., Salmond, S. J. and Smith, A. F. M. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings-F* **140** 107–113.
[25] Huggins, J. H. and Zou, J. (2016). Quantifying the accuracy of approximate diffusions and Markov chains. *ArXiv e-prints.*
[26] Jacob, P. E. and Thiery, A. H. (2015). On non-negative unbiased estimators. *Annals of Statistics* **43** 769–784.
[27] Jacod, J. and Protter, P. (2012). *Discretization of Processes.*, 1st ed. Springer, Berlin.
[28] Johansen, A. M. and Doucet, A. (2008). A Note on the Auxiliary Particle Filter. *Statistics and Probability Letters* **78** 1498–1504.
[29] Johnson, R. A. (1970). Asymptotic expansions associated with posterior distributions. *The Annals of Mathematical Statistics* **41** 851–864.
[30] Jordan, M. I. (2013). On statistics, computation and scalability. *Bernoulli* **19** 1378–1390.
[31] Karatzas, I. and Shreve, S. (1991). *Brownian Motion and Stochastic Calculus.*, 2nd ed. Springer, New York.
[32] Kingman, J. F. C. (1992). *Poisson Processes.*, 1st ed. Clarendon Press,

Oxford.

[33] Kloeden, P. E. and Platen, E. (1992). *Numerical Solution of Stochastic Differential Equations.*, 4th ed. Springer, Berlin.

[34] Kong, A., Liu, J. S. and Wong, W. H. (1994). Sequential Imputations and Bayesian Missing Data Problems. *Journal of the American Statistical Association* **89** 278–288.

[35] Korattikara, A., Chen, Y. and Welling, M. (2013). Austerity in MCMC land: Cutting the Metropolis-Hastings budget. *arXiv preprint arXiv:1304.5299*.

[36] Letac, G. (1975). On building random variables of a given distribution. *The Annals of Probability* **3** 298–306.

[37] Li, C., Srivastava, S. and Dunson, D. B. (2016). Simple, Scalable and Accurate Posterior Interval Estimation. *arXiv preprint arXiv:1605.04029*.

[38] Ma, Y., Chen, T. and Fox, E. B. (2015). A complete recipe for stochastic gradient MCMC. In *Advances in Neural Information Processing Systems* 2917–2925.

[39] Milicer, H. and Szczotka, F. (1966). Age at Menarche in Warsaw girls in 1965. *Human Biology* **38** 199-203.

[40] Minsker, S., Srivastava, S., Lin, L. and Dunson, D. B. (2014). Robust and Scalable Bayes via a Median of Subset Posterior Measures Mathematics e-print No. 1403.2660, ArXiv.

[41] Del Moral, P. (2004). *Feynman-Kac formulae: genealogical and interacting particle systems with applications. Probability and Its Applications.* Springer Verlag, New York.

[42] Del Moral, P., Doucet, A. and Jasra, A. (2012). On Adaptive Resampling Procedures for Sequential Monte Carlo Methods. *Bernoulli* **18** 252–278.

[43] Neiswanger, W., Wang, C. and Xing, E. (2013). Asymptotically Exact, Embarrassingly Parallel MCMC Mathematics e-print No. 1403.2660, ArXiv.

[44] Øksendal, B. (2007). *Stochastic Differential Equations.*, 6th ed. Springer, Berlin.

[45] Platen, E. and Bruti-Liberati, N. (2010). *Numerical Solution of Stochastic Differential Equations with Jumps in Finance.*, 1st ed. Springer, Berlin.

[46] Pollock, M. (2013). Some Monte Carlo Methods for Jump Diffusions. PhD thesis, Department of Statistics, University of Warwick.

[47] Pollock, M. (2015). On the Exact Simulation of (Jump) Diffusion Bridges. *Proceedings of the 2015 Winter Simulation Conference.*

[48] Pollock, M., Johansen, A. M. and Roberts, G. O. (2016). On the Exact and ϵ-Strong Simulation of (Jump) Diffusions. *Bernoulli* **22** 794–856.

[49] Robert, C. and Casella, G. (2004). *Monte Carlo statistical methods*, second ed. Springer Science & Business Media.

[50] Roberts, G. O. and Tweedie, R. L. (1996). Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli* **2** 341–363.

[51] Scott, S. L., Blocker, A. W. and Bonassi, F. V. (2016). Bayes and Big Data: The Consensus Monte Carlo Algorithm. *International Journal of Management Science and Engineering Management*.
[52] Srivastava, S., Cevher, V., Tran-Dinh, Q. and Dunson, D. B. (2015). WASP: Scalable Bayes via barycenters of subset posteriors. In *AISTATS*.
[53] Steinsaltz, D. and Evans, S. (2007). Quasistationary distributions for one-dimensional diffusions with killing. *Transactions of the American Mathematical Society* **359** 1285–1324.
[54] Tanner, M. A. and Wong, W. H. (1987). The calculation of posterior distributions by data augmentation. *Journal of the American statistical Association* **82** 528–540.
[55] Taylor, K. B. (2015). Exact Algorithms for simulation of diffusions with discontinuous drift and robust Curvature Metropolis-adjusted Langevin algorithms. PhD thesis, Department of Statistics, University of Warwick.
[56] R Core Team (2013). R: A Language and Environment for Statistical Computing R Foundation for Statistical Computing, Vienna, Austria.
[57] Venables, W. N. and Ripley, B. D. (2002). *Modern Applied Statistics with S*. Springer.
[58] Vollmer, S. J., Zygalakis, K. C. and Teh, Y. W. (2015). (Non-) asymptotic properties of Stochastic Gradient Langevin Dynamics. *ArXiv e-prints*.
[59] Wang, X. and Dunson, D. B. (2013). Parallelizing MCMC via Weierstrass Sampler Mathematics e-print No. 1312.4605, ArXiv.
[60] Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient Langevin dynamics. *Proceedings of the 28th International Conference on Machine Learning (ICML-11)* 681–688.
[61] W., T. Y., Thiéry, A. H. and Vollmer, S. J. (2016). Consistency and fluctuations for stochastic gradient Langevin dynamics. *Journal of Machine Learning Research* **17** 1–33.

Appendix A: Polynomial tails

In this appendix we examine the condition in (14) we use within Theorem 1. This is essentially a condition on the tail of ν, and so we examine the critical case in which the tails of ν are heavy. More precisely, we demonstrate that for polynomial tailed densities in one-dimension that (14) essentially amounts to requiring that $\nu^{1/2}$ is integrable. Recall that by construction $\nu^{1/2}$ will be integrable as we have chosen $\nu^{1/2} = \pi$.

For simplicity, suppose that ν is a density on $[1, \infty)$ such that $\nu(x) = x^{-p}$. In this case we can easily compute that for $p > 1$,

$$\nabla\nu(x) = -px^{-p-1}$$
$$\Delta\nu(x) = p(p+1)x^{-p-2}$$

from which we can easily compute the quantity whose limit is taken in (14) as

$$x^{p(\gamma-1/2)-2}[p(p+1)-\gamma p^2].$$

As such, we have that condition (14) holds if and only if

$$p+1-\gamma p>0 \tag{44}$$

and

$$p(\gamma-1/2)-2\geq 0. \tag{45}$$

Now we shall demonstrate that we can find such γ for all $p>2$. For instance, suppose that $p=2+\epsilon$. The case $\epsilon\geq 2$ can be handled by just setting $\gamma=1$, so suppose otherwise and set $\gamma=3/2-\epsilon/4$. In this case, (45) just gives $\epsilon/2-\epsilon^2/4\geq 0$. Moreover the expression in (44) becomes $3\epsilon/2+\epsilon^2>0$, completing our argument.

Appendix B: Conditions for constructing Path-space Rejection Samplers

In [48] a framework for constructing path-space rejection samplers (for diffusions and jump diffusions) was established in which entire diffusion sample paths could be represented by means of simulating a finite dimensional *skeleton* (which we generalise here for d-dimensional diffusions with non-unit volatility matrix, such as those in (7)). A sample path skeleton is comprised of a *layer* constraining the sample path, and is sufficient to recover the sample path at any other finite collection of time points without error as desired. In this section, we will provide only a brief overview of the key notions and conditions, which the reader should note have been weakened in order to ease exposition for this particular paper — for a fuller exposition please refer to [46, 48, 47].

Definition 1 (Skeleton, [46])**.** *A skeleton ($\mathcal{S}$) is a finite dimensional representation of a diffusion sample path ($\boldsymbol{X}\sim\mathbb{Q}^{\mathbf{x}}$), that can be simulated without any approximation error by means of a proposal sample path drawn from an equivalent proposal measure ($\mathbb{W}^{\mathbf{x}}$) and accepted with probability proportional to $\frac{\mathrm{d}\mathbb{Q}^{\mathbf{x}}}{\mathrm{d}\mathbb{W}^{\mathbf{x}}}(\mathbf{X})$, which is sufficient to restore the sample path at any finite collection of time points exactly with finite computation where $\mathbf{X}|\mathcal{S}\sim\mathbb{W}^{\mathbf{x}}|\mathcal{S}$.*

Definition 2 (Layer, [46])**.** *A layer $R(\mathbf{X})$, is a function of a diffusion sample path $\mathbf{X}\sim\mathbb{W}^{\mathbf{x}}$ which determines the compact hypercuboid to which any particular sample path $\mathbf{X}(\omega)$ is constrained.*

Clearly, as we are ultimately constructing a method which relies upon every step being exact then we need to ensure that we can determine exactly which layer to ascribe a given (partially realised) sample path ($R(\mathbf{X})\sim\mathcal{R}:=\mathbb{W}^{\mathbf{x}}\circ R^{-1}$), and that we can simulate any finite collection of intermediate points of the trajectory conditional on the skeleton $\mathbf{X}_{t_1},\ldots,\mathbf{X}_{t_n}\sim\mathbb{W}^{\mathbf{x}}|\mathcal{S}$ (which recall includes

the layer).

Further defining the function $B : \mathbb{R}^d \to \mathbb{R}$ such that by $\beta(\mathbf{u}) := \nabla B(\mathbf{u})$ and setting $\phi(\mathbf{u}) := (||\boldsymbol{\Lambda}^{-1/2}\beta(\mathbf{u})||^2 + \operatorname{div}\boldsymbol{\Lambda}^{-1/2}\beta(\mathbf{u}))/2$, then typically conditions similar to the following conditions (1–3) are imposed on PRS [46],

Condition 1 (Solutions)**.** *The coefficients of (7) are sufficiently regular to ensure the existence of a unique, non-explosive, weak solution.*

Condition 2 (Continuity)**.** *The drift coefficient $\beta_i \in C^1$ $\forall i \in \{1, \dots, d\}$.*

Condition 3 (Growth Bound)**.** *We have that $\exists K > 0$ such that $|\beta_i(\mathbf{u})| \leq K(1 + ||\mathbf{u}||)$ $\forall \mathbf{u} \in \mathbb{R}^d$ and $i \in \{1, \dots, d\}$.*

however, recalling that we are interested in drifts of the form $\beta(\mathbf{u}) := \frac{1}{2}\boldsymbol{\Lambda}\nabla \log \nu(\mathbf{u})$ and that (by construction) we have that (7) has invariant measure ν, then Condition 1 is satisfied. Conditions 2–3 will also be satisfied for a very broad class of interesting problems, and could be substantively weakened by direct application of [55] (however this adds complication that detracts from the primary results of this paper). Imposition of Condition 2 is necessary to ensure Result 2 holds, and Conditions 2 and 3 are necessary for establishing Result 1 (both of which follow).

Now, as indicated in Definition 1, in order to deploy PRS we need to establish that the Radon-Nikodým derivative of $\mathbb{Q}^{\mathbf{x}}$ with respect to $\mathbb{W}^{\mathbf{x}}$ exists (Results 1) and can be bounded (above and below) on compact sets (Result 2). In particular, we have,

Result 1 (Unconditioned Radon-Nikodým derivative)**.** *As noted in [44, 8], under Conditions 1–3, the Radon-Nikodým derivative of $\mathbb{Q}^{\mathbf{x}}$ with respect to $\mathbb{W}^{\mathbf{x}}$ exists and is given by Girsanov's formula as follows,*

$$\frac{\mathrm{d}\mathbb{Q}^{\mathbf{x}}}{\mathrm{d}\mathbb{W}^{\mathbf{x}}}(\mathbf{X}) = \exp\left\{B(\mathbf{X}_T) - B(\mathbf{x}) - \int_0^T \phi(\mathbf{X}_s)\,\mathrm{d}s\right\}. \tag{46}$$

Corollary 1 (Unconditioned Radon-Nikodým derivative: Langevin Case)**.** *Recalling that in the case of the Langevin diffusion $\beta(\mathbf{x}) := \frac{1}{2}\boldsymbol{\Lambda}\nabla \log \nu(\mathbf{x})$, by direct application of Result 1 we have,*

$$\frac{\mathrm{d}\mathbb{Q}^{\mathbf{x}}}{\mathrm{d}\mathbb{W}^{\mathbf{x}}}(\mathbf{X}) \propto \nu^{1/2}(\mathbf{X}_T) \cdot \exp\left\{-\int_0^T \phi(\mathbf{X}_s)\,\mathrm{d}s\right\}. \tag{47}$$

Now, as the algorithm is fundamentally a rejection sampler, and as noted in [36] and [18, Sec. II 3.5] it is necessary to have a constant bound on the Radon-Nikodým derivative, then we impose the following two additional conditions on the Radon-Nikodým derivative of Result 1.

Condition 4 (Φ)**.** *There exists a constant $\Phi > -\infty$ such that $\Phi \leq \phi(\mathbf{u})$ $\forall \mathbf{u} \in \mathbb{R}^d$.*

Condition 5 (Ψ)**.** *There exists a constant $\Psi < \infty$ such that $B(\mathbf{u}) \leq \Psi$ $\forall \mathbf{u} \in \mathbb{R}^d$.*

In order to circumvent imposing further conditions on β (in order to explicitly bound ϕ and $\frac{\mathrm{d}\mathbb{Q}^{\mathbf{x}}}{\mathrm{d}\mathbb{W}^{\mathbf{x}}}(\mathbf{X})$ above and below), this paper exploits a localised construction (using the *layer* of Definition 2) of draws from the proposal measure ($\mathbf{X} \sim \mathbb{W}^{\mathbf{x}}$), which is sufficient due to the following result,

Result 2 (Local Boundedness, [46])**.** *Condition 2 establishes that we have β_i and $(div\ \beta)_i$ bounded on compact sets $\forall i \in \{1, \ldots, d\}$. As such, we have that $\exists\, \ell_i, \upsilon_i \in \mathbb{R}$ such that $\forall\ t \in [0,T]$, $\mathbf{X}_t(\omega) \in [\ell_1, \upsilon_1] \times \ldots \times [\ell_d, \upsilon_d]$ and so $\exists\, L_{\mathbf{X}} := L\left(\mathbf{X}(\omega)\right) \in \mathbb{R}, U_{\mathbf{X}} := U\left(\mathbf{X}(\omega)\right) \in \mathbb{R}$ such that $\forall\ t \in [0,T]$, $\phi\left(\mathbf{X}_t(\omega)\right) \in [L_{\mathbf{X}}, U_{\mathbf{X}}]$.*

As we have now established the existence and boundedness of the Radon-Nikodým derivative of interest, then we can find the explicit representation of the transition density of (7) which is used in the main body of the paper. In particular, denoting by $\mathbb{Q}^{\mathbf{x},\mathbf{y}}$ and $\mathbb{W}^{\mathbf{x},\mathbf{y}}$ as the measure induced by (7) and the driftless version of (7) with the additional constraint that $\mathbf{X}_T = \mathbf{y} \in \mathbb{R}$, $p_T(\mathbf{x},\mathbf{y}) := \mathbb{P}_{\mathbb{Q}^{\mathbf{x}}}(\mathbf{X}_T \in \mathrm{d}\mathbf{y} \,|\, \mathbf{X}_0 = \mathbf{x})/\,\mathrm{d}\mathbf{y}$ and $w_T(\mathbf{x},\mathbf{y}) := \mathbb{P}_{\mathbb{W}^{\mathbf{x}}}(\mathbf{X}_T \in \mathrm{d}\mathbf{y} \,|\, \mathbf{X}_0 = \mathbf{x})/\,\mathrm{d}\mathbf{y}$ as the transition densities (7) and the driftless version of (7), and considering the Radon-Nikodým derivative of $\mathbb{Q}^{\mathbf{x},\mathbf{y}}$ with respect to $\mathbb{W}^{\mathbf{x},\mathbf{y}}$, we have,

Result 3 (Conditioned Radon-Nikodým derivative)**.** *Following directly from Result 1 we have [17],*

$$\frac{\mathrm{d}\mathbb{Q}^{\mathbf{x},\mathbf{y}}}{\mathrm{d}\mathbb{W}^{\mathbf{x},\mathbf{y}}}(\mathbf{X}) = \frac{w_T(\mathbf{x},\mathbf{y})}{p_T(\mathbf{x},\mathbf{y})} \cdot \frac{\mathrm{d}\mathbb{Q}^{\mathbf{x}}}{\mathrm{d}\mathbb{W}^{\mathbf{x}}}(\mathbf{X}),$$

with transition density of the following form (by taking expectations with respect to $\mathbb{W}^{\mathbf{x},\mathbf{y}}$),

$$p_T(\mathbf{x},\mathbf{y}) = w_T(\mathbf{x},\mathbf{y}) \cdot \mathbb{E}_{\mathbb{W}^{\mathbf{x},\mathbf{y}}}\left[\frac{\mathrm{d}\mathbb{Q}^{\mathbf{x}}}{\mathrm{d}\mathbb{W}^{\mathbf{x}}}(\mathbf{X})\right]. \tag{48}$$

Corollary 2 (Conditioned Radon-Nikodým derivative: Langevin Case)**.** *Once again, recalling that in the case of the Langevin diffusion $\beta(\mathbf{x}) := \frac{1}{2}\boldsymbol{\Lambda}\nabla \log \nu(\mathbf{x})$, then by direct application of Result 3 we have,*

$$p_T(\mathbf{x},\mathbf{y}) \propto w_T(\mathbf{x},\mathbf{y}) \cdot \nu^{1/2}(\mathbf{y}) \cdot \mathbb{E}_{\mathbb{W}^{\mathbf{x},\mathbf{y}}}\left[\exp\left\{-\int_0^T \phi(\mathbf{X}_s)\,\mathrm{d}s\right\}\right]. \tag{49}$$

Appendix C: Simulation of a Path-Space Layer and Intermediate Points

In this appendix we present the methodology and algorithms required for simulating an individual proposal trajectory of (layered) killed multivariate Brownian motion, which is what is required in Section 3. Our exposition is as follows: In Appendix C.1 we present the work of [19], in which a highly efficient rejection sampler is developed (based on the earlier work of [11]) for simulating the first

passage time for univariate standard Brownian motion for a given symmetric boundary, extending it to consider the case of the univariate first passage times of d-dimensional standard Brownian motion with non-symmetric boundaries. This construction allows us to determine an interval (given by the first, first passage time) and layer (a hypercuboid inscribed by the user specified univariate boundaries) in which the sample path is almost-surely constrained, and by application of the strong Markov property can be applied iteratively to find for any interval of time a layer (a concatenation of hypercuboids) which almost-surely constrains the sample path; In Appendix C.2 we present a rejection sampler enabling the simulation of constrained univariate standard Brownian motion as developed in Section C.1, at any desired intermediate point. As motivated in Section 3 these intermediate points may be at some random time (corresponding to a proposed killing point of the proposed sample path), or a deterministic time (in which the sample path is extracted for inclusion within the desired Monte Carlo estimator of QSMC (29)); Finally, in Appendix C.3 we present the full methodology required in Sections 3 and 4 in which we simulate multivariate Brownian motion at any desired time marginal, with d-dimensional hypercuboids inscribing intervals of the state space in which the sample path almost surely lies.

C.1. Simulating the first passage times of univariate and multivariate standard Brownian motion

To begin with we restrict our attention to the (i^{th}) dimension of multivariate standard Brownian motion initialised at 0, and the first passage time of the level $\theta^{(i)}$ (which is specified by the user). In particular we denote,

$$\tau^{(i)} := \inf\{t \in \mathbb{R}+ \,:\, |W_t^{(i)} - W_0^{(i)}| \geq \theta^{(i)}\}. \tag{50}$$

Recalling the self similarity properties of Brownian motion ([31, Section 2.9]), we can further restrict our attention to the simulation of the first passage time of univariate Brownian motion of the level 1, noting that $\tau^{(i)} \overset{\mathcal{D}}{=} \left(\theta^{(i)}\right)^2 \bar{\tau}$ where,

$$\bar{\tau} := \inf\{t \in \mathbb{R}+ \,:\, |W_t - W_0| \geq 1\}, \tag{51}$$

noting that at this level,

$$\mathbb{P}(W_\tau = W_0 + 1) = \mathbb{P}(W_\tau = W_0 - 1) = \frac{1}{2}. \tag{52}$$

Denoting by $f_{\bar{\tau}}$ the density of $\bar{\tau}$ (which cannot be evaluated point-wise), then the approach outlined in [19] for drawing random samples from $f_{\bar{\tau}}$ is a series sampler. In particular, an accessible dominating density of $f_{\bar{\tau}}$ is found (denoted $g_{\bar{\tau}}$) from which exact proposals can be made, then upper and lower monotonically convergent bounding functions are constructed ($\lim_{n\to\infty} f^{\uparrow}_{\bar{\tau},n} \to f_{\bar{\tau}}$ and $\lim_{n\to\infty} f^{\downarrow}_{\bar{\tau},n} \to f_{\bar{\tau}}$ such that for any $t \in \mathbb{R}_+$ and $\epsilon > 0$ $\exists\, n^*(t)$ such that

$\forall\, n \geq n^*(t)$ we have $f^{\uparrow}_{\bar{\tau},n}(t) - f^{\downarrow}_{\bar{\tau},n}(t) < \epsilon$), and then evaluated to sufficient precision such that acceptance or rejection can be made while retaining exactness. A minor complication arises in that no single dominating density is uniformly efficient on $\mathbb{R}_+$, and furthermore no single representation of the bounding functions monotonically converge to the target density point-wise on $\mathbb{R}_+$. As such, the strategy deployed by [19] is to exploit a dual representation of $f_{\bar{\tau}}$ given by [15] in order to construct a hybrid series sampler, using one representation of $f_{\bar{\tau}}$ for the construction of a series sampler on the interval $(0, t_1]$ and the other representation for the interval $[t_2, \infty)$ (fortunately we have $t_1 > t_2$, and so we have freedom to choose a threshold $t^* \in [t_2, t_1]$ in which to splice the series samplers). In particular, as shown in [15] $f_{\bar{\tau}}(t) = \pi \sum_{k=0}^{\infty} (-1)^k a_k(t)$ where,

$$a_k(t) = \begin{cases} \left(\frac{2}{\pi t}\right)^{3/2} \left(k + \frac{1}{2}\right) \exp\left\{-\frac{2}{t}(k+\frac{1}{2})^2\right\}, & (1) \\ \left(k + \frac{1}{2}\right) \exp\left\{-\frac{1}{2}(k+\frac{1}{2})^2 \pi^2 t\right\}, & (2) \end{cases} \tag{53}$$

and so by consequence upper and lower bounding sequences can be constructed by simply taking either representation and truncating the infinite sum to have an odd or even number of terms respectively. More precisely,

$$f^{\downarrow}_{\bar{\tau},n}(t) := \left(\pi \sum_{k=0}^{2n+1} (-1)^k a_k(t)\right)_+, \qquad f^{\uparrow}_{\bar{\tau},n}(t) := \left[\pi \sum_{k=0}^{2n+1} (-1)^k a_k(t)\right] \wedge g_{\bar{\tau}}(t). \tag{54}$$

As shown in [19, Lemma 1], the bounding sequences based on the representation of $f_{\bar{\tau}}(t)$ in (53.1) are monotonically converging for $t \in (0, 4/\log(3)]$, and for (53.2) monotonically converging for $t \in [\log(3)/\pi^2, \infty)$. After choosing a suitable threshold $t^* \in [4/\log(3), \log(3)/\pi^2]$ for which to splice the series samplers, then by simply taking the first term in each representation of $f_{\bar{\tau}}(t)$ a dominating density can be constructed as follows,

$$f_{\bar{\tau}}(t) \leq g_{\bar{\tau}}(t) \propto \underbrace{\frac{2}{\pi t^{3/2}} \exp\left\{-\frac{1}{2t}\right\} \cdot \mathbb{1}_{t \leq t^*}}_{\propto g^{(1)}_{\bar{\tau}}(t)} + \underbrace{\frac{\pi}{2} \exp\left\{-\frac{\pi^2 t}{8}\right\} \cdot \mathbb{1}_{t \geq t^*}}_{\propto g^{(2)}_{\bar{\tau}}(t)}. \tag{55}$$

[19] empirically optimises the choice of $t^* = 0.64$ so as to minimise the normalising constant of (55). With this choice $M_1 := \int g^{(1)}_{\bar{\tau}}(t)\, \mathrm{d}t \approx 0.422599$ (to 6 d.p.) and $M_2 := \int g^{(2)}_{\bar{\tau}}(t)\, \mathrm{d}t \approx 0.578103$ (to 6 d.p.), and so we have a normalising constant $M = M_1 + M_2 \approx 1.000702$ (to 6 d.p.) which equates to the expected number of proposal random samples drawn from $g_{\bar{\tau}}$ before one would expect an accepted draw (the algorithmic 'outer loop'). Now considering the iterative algorithmic 'inner loop' – in which the bounding sequences are evaluated to precision sufficient to determine acceptance or rejection – as shown in [19], the exponential convergence of the sequences ensures that in expectation this is uniformly bounded by 3.

Simulation from $g_{\bar{\tau}}$ is possible by either simulating $\bar{\tau} \sim g^{(1)}_{\bar{\tau}}$ with probability M_1/M, else $\bar{\tau} \sim g^{(2)}_{\bar{\tau}}$. Simulating $\bar{\tau} \sim g^{(1)}_{\bar{\tau}}$ can be achieved by noting $t \overset{\mathcal{D}}{=} t^* + 8X/\pi^2$, where $X \sim \text{Exp}(1)$. Simulating $\bar{\tau} \sim g^{(2)}_{\bar{\tau}}$ can be achieved by noting that as outlined in [18, IX.1.2] $t \overset{\mathcal{D}}{=} t^*/(1+t^*X)^2$, where $X := \inf_i\{\{X_i\}_{i=1}^{\infty} \overset{\text{iid}}{\sim} \text{Exp}(1) : (X_i)^2 \leq 2X_{i+1}/t^*, (i-1)/2 \in \mathbb{Z}\}$.

A summary of the above for simulating jointly the first passage time and location of the i^{th} dimension of Brownian motion of the threshold level $\theta^{(i)}$ is provided in Algorithm 4.

Note that generalising to the case where we are interested in the first pas-

Algorithm 4 Simulating $(\tau, W^{(i)}_{\tau})$, where $\tau := \inf\{t \in \mathbb{R}+ : |W^{(i)}_t - W^{(i)}_0| \geq \theta^{(i)}\}$ [19].

1. Input $W^{(i)}_0$ and $\theta^{(i)}$.
2. $g_{\bar{\tau}}$: Simulate $u \sim \text{U}[0,1]$,
   (a) $g^{(1)}_{\bar{\tau}}$: If $u \leq M_1/M$, then simulate $X \sim \text{Exp}(1)$ and set $\bar{\tau} := t^* + 8X/\pi^2$.
   (b) $g^{(2)}_{\bar{\tau}}$: If $u > M_1/M$, then set $X := \inf_i\{\{X_i\}_{i=1}^{\infty} \overset{\text{iid}}{\sim} \text{Exp}(1) : (X_i)^2 \leq 2X_{i+1}/t^*, (i-1)/2 \in \mathbb{Z}\}$ and set $\bar{\tau} := t^*/(1+t^*X)^2$.
3. u: Simulate $u \sim \text{U}[0,1]$ and set $n = 0$.
4. $f_{\bar{\tau},n}$: While $u \cdot g_{\bar{\tau}}(\bar{\tau}) \in (f^{\downarrow}_{\bar{\tau},n}(\bar{\tau}), f^{\uparrow}_{\bar{\tau},n}(\bar{\tau}))$, set $n = n+1$.
5. $f_{\bar{\tau}}$: If $u \cdot g_{\bar{\tau}}(\bar{\tau}) \leq f^{\downarrow}_{\bar{\tau},n}(\bar{\tau})$ accept, else reject and return to Step 2.
6. τ: Set $\tau := (\theta^{(i)})^2\bar{\tau}$.
7. $W^{(i)}_{\tau}$: With probability 1/2 set $W^{(i)}_{\tau} = W^{(i)}_0 + \theta^{(i)}$, else set $W^{(i)}_{\tau} = W^{(i)}_0 - \theta^{(i)}$.
8. Return $(\tau, W^{(i)}_{\tau})$.

sage time of Brownian motion of a non-symmetric barrier, in particular for $\ell^{(i)}, \upsilon^{(i)} \in \mathbb{R}_+$,

$$\tau^{(i)} := \inf\{t \in \mathbb{R}+ : W^{(i)}_t - W^{(i)}_0 \notin (W^{(i)}_0 - \ell^{(i)}, W^{(i)}_0 + \upsilon^{(i)})\}, \tag{56}$$

is trivial algorithmically. In particular, using the strong Markov property we can iteratively apply Algorithm 4 setting $\theta^{(i)} := \min(\ell^{(i)}, \upsilon^{(i)})$ and simulating intermediate first passage times of lesser barriers, halting whenever the desired barrier is attained. We suppress this (desirable) flexibility in the remainder of the paper to avoid the resulting notational complexity.

C.2. Simulating intermediate points of multivariate standard Brownian motion conditioned on univariate first passage times

Clearly in addition to being able to simulate the first passage times of a single dimension of Brownian motion, we want to be able simulate the remainder of the dimensions of Brownian motion at that time, or indeed the sample path at times other than its first passage times. As the dimensions of Brownian motion are independent (and so Brownian motion can be composed by considering each dimension separately), we can restrict our attention to simulating a single dimension of the sample path for an intermediate time $q \in [s, \tau]$ given W_s, the extremal value W_τ, and constrained such that $\forall u \in [s, \tau], W_u \in [W_s - \theta, W_s + \theta]$. Furthermore, as we are only interested in the forward simulation of Brownian motion (as motivated in Section 3.1), then by application of the strong Markov property we need only consider the simulation of a single intermediate point (although note by application of [48, Section 7] simulation at times conditional on future information is possible).

To proceed, note that (as outlined in [2, Prop. 2]) the law of a univariate Brownian motion sample path in the interval $[s, \tau]$ (where $s < \tau$) initialised at (s, W_s) and constrained to attain its extremal value at (τ, W_τ), is simply the law of a three dimensional Bessel bridge. We require the additional constraint that $\forall u \in [s, \tau], W_u \in [W_s - \theta, W_s + \theta]$, which can be imposed in simulation by deploying a rejection sampling scheme in which a Bessel bridge sample path is simulated at a single required point (as above) and accepted if it meets the imposed constraint at either side of the simulated point, and rejected otherwise.

As presented in [6, 46], the law of a Bessel bridge sample path (parametrised as above) coincides with that of an appropriate time rescaling of three independent Brownian bridge sample paths of unit length conditioned to start and end at the origin (denoted by $\{b^{(i)}\}_{i=1}^3$). Supposing we require the realisation of a Bessel bridge sample path at some time $q \in [s, \tau]$, then by simply realising three independent Brownian bridge sample paths at that time marginal ($\{b_q^{(i)}\}_{i=1}^3$), we have,

$$W_q = W_s + (-1)^{\mathbb{1}(W_\tau < W_s)} \sqrt{(\tau - s)\left[\left(\frac{\theta(\tau - q)}{(\tau - s)^{3/2}} + b_q^{(1)}\right)^2 + (b_q^{(2)})^2 + (b_q^{(3)})^2\right]}. \tag{57}$$

The method by which the proposed Bessel bridge intermediate point is accepted or rejected (recall, to impose the constraint that $\forall u \in [s, \tau], W_u \in [W_s - \theta, W_s + \theta]$) is non-trivial as there does not exist a closed form representation of the required probability (which we will denote in this appendix by p). Instead, as shown in Theorem 3, a representation for p can be found as the product of two infinite series, which as a consequence of this form can not be evaluated directly in order to make the typical acceptance-rejection comparison (i.e. determining whether $u \leq p$ or $u > p$, where $u \sim \mathrm{U}[0, 1]$). The

strategy we deploy to retain exactness and accept with the correct probability p is that of a retrospective Bernoulli sampler [48, Sec. 6.0]. In particular, in Corollary 3 we construct monotonically convergent upper and lower bounding probabilities ($p_n^{\uparrow}$ and $p_n^{\downarrow}$ respectively) with the property that $\lim_{n\to\infty} p_n^{\uparrow} \to p$ and $\lim_{n\to\infty} p_n^{\downarrow} \to p$ such that for any $u \in [0,1]$ and $\epsilon > 0$ $\exists\, n^*(t)$ such that $\forall\, n \geq n^*(t)$ we have $p_n^{\uparrow} - p_n^{\downarrow} < \epsilon$, which are then evaluated to sufficient precision to make the acceptance-rejection decision, taking almost surely finite computational time.

Theorem 3. *The probability that a three dimensional Bessel bridge sample path $W \sim \mathbb{W}_{s,\tau}^{W_s,W_\tau} \,\big|\, (W_\tau, W_q)$ for $s < q < \tau$ attaining its boundary value at (τ, W_τ), remains in the interval $[W_s - \theta, W_s + \theta]$, can be represented by the following product of infinite series (where we denote by $m := \mathbb{1}(W_\tau > W_s) - \mathbb{1}(W_\tau < W_s)$),*

$$\begin{aligned}\mathbb{P}\left(W_{[s,\tau]} \in [W_s-\theta, W_s+\theta] | W_s, W_q, W_\tau\right) &= \underbrace{\left(\frac{1 - \sum_{j=1}^{\infty}\left[\varsigma_{q-s}(j; W_s - W_q, \theta) - \varphi_{q-s}(j; W_s - W_q, \theta)\right]}{1 - \exp\left\{-2\theta[m(W_s - W_q) + \theta]/(q-s)\right\}}\right)}_{=:p_1} \\ &\quad \cdot \underbrace{\left(1 + \sum_{j=1}^{\infty}\left[\psi_{\tau-q}(j; W_q - W_\tau, \theta, m) + \chi_{\tau-q}(j; W_q - W_\tau, \theta, m)\right]\right)}_{=:p_2},\end{aligned} \tag{58}$$

where,

$$\varsigma_{\Delta}(j; \delta, \theta) := 2 \cdot \exp\left\{-\frac{2\theta^2(2j-1)^2}{\Delta}\right\} \cdot \cosh\left(\frac{2(2j-1)\theta\delta}{\Delta}\right), \tag{59}$$

$$\varphi_{\Delta}(j; \delta, \theta) := 2 \cdot \exp\left\{-\frac{8\theta^2 j^2}{\Delta}\right\} \cdot \cosh\left\{\frac{4\theta\delta j}{\Delta}\right\}, \tag{60}$$

$$\psi_{\Delta}(j; \delta, \theta, m) := \chi_{\Delta}(j; \delta, \theta, -m) := \frac{(4\theta j + m\delta)}{m\delta} \cdot \exp\left\{-\frac{4\theta j}{\Delta}(2\theta j + m\delta)\right\}. \tag{61}$$

Proof. Begin by noting that that the strong Markov property allows us to decompose our required probability as follows,

$$\begin{aligned}&\mathbb{P}\left(W_{[s,\tau]} \in [W_s - \theta, W_s + \theta] | W_s, W_q, W_\tau\right) \\ &= \underbrace{\mathbb{P}\left(W_{[s,q]} \in [W_s - \theta, W_s + \theta] | W_s, W_q\right)}_{p_1} \cdot \underbrace{\mathbb{P}\left(W_{[q,\tau]} \in [W_s - \theta, W_s + \theta] | W_q, W_\tau\right)}_{p_2}.\end{aligned} \tag{62}$$

Relating the decomposition to the statement of the theorem, p_1 follows directly from the parametrisation given and the representation in [46, Thm. 6.1.2] of the result in [7, Prop. 3]. p_2 similarly follows from the representation found in [48, Thm. 5].

Corollary 3. *Letting* $p := \mathbb{P}\left(W_{[s,\tau]} \in [W_s-\theta, W_s+\theta]\right)$*, monotonically convergent upper and lower bounding probabilities* ($p_n^{\uparrow}$ *and* $p_n^{\downarrow}$ *respectively) with the property that* $\lim_{n\to\infty} p_n^{\uparrow} \to p$ *and* $\lim_{n\to\infty} p_n^{\downarrow} \to p$ *can be found (where* $n_0 := \lceil \sqrt{(\tau-q)+4\theta^2}/4\theta \rceil$*),*

$$p_n^{\downarrow} := \left(\frac{1 - \sum_{j=1}^{n} \varsigma_{q-s}(j; W_s - W_q, \theta) + \sum_{j=1}^{n-1} \varphi_{q-s}(j; W_s - W_q, \theta)}{1 - \exp\left\{-2\theta[m(W_s - W_q) + \theta]/(q-s)\right\}} \right) \cdot \left(1 + \sum_{j=1}^{n_0+n} \psi_{\tau-q}(j; W_q - W_\tau, \theta, m) + \sum_{j=1}^{n_0+n-1} \chi_{\tau-q}(j; W_q - W_\tau, \theta, m)\right], \tag{63}$$

$$p_n^{\uparrow} := \left(\frac{1 - \sum_{j=1}^{n} \varsigma_{q-s}(j; W_s - W_q, \theta) + \sum_{j=1}^{n} \varphi_{q-s}(j; W_s - W_q, \theta)}{1 - \exp\left\{-2\theta[m(W_s - W_q) + \theta]/(q-s)\right\}} \right) \cdot \left(1 + \sum_{j=1}^{n_0+n} \psi_{\tau-q}(j; W_q - W_\tau, \theta, m) + \sum_{j=1}^{n_0+n} \chi_{\tau-q}(j; W_q - W_\tau, \theta, m)\right]. \tag{64}$$

Furthermore we have

$$\frac{p_n^{\uparrow} - p_n^{\downarrow}}{p_{n-1}^{\uparrow} - p_{n-1}^{\downarrow}} =: r_n \leq r \in (0,1), \tag{65}$$

and so,

$$\bar{K} := \sum_{i=1}^{\infty} |p_i^{\uparrow} - p_i^{\downarrow}| = (p_1^{\uparrow} - p_1^{\downarrow}) + \sum_{i=2}^{\infty} \prod_{j=2}^{i} r_j \leq \sum_{i=0}^{\infty} r^i = \frac{1}{1-r} < \infty. \tag{66}$$

Proof. The summations in the left hand brackets of the sequences (63) and (64) follows from Theorem 3 and [7, Prop. 3]. The summations in the right hand brackets of the sequences (63) and (64), and the necessary condition on n_0, follows from [48, Corollary 5]. The validity of the product form of (63) and (64) follows from [48, Corollary 1]. The bound on the ratio of subsequent bound ranges of p in (65) follows from the exponential decay in n of $\varsigma(n)$, $\varphi(n)$, $\psi(n)$ and $\chi(n)$ of Theorem 3, and as shown in the proof of [46, Thm. 6.1.1] and [46, Corollary 6.1.3]. (66) follows directly from (65). □

Having established Theorem 3 and Corollary 3 we can now construct a (retrospective) rejection sampler in which we simulate W_q (as per the law of a Bessel

bridge) and, by means of an algorithmic loop in which the bounding sequences of the acceptance probability are evaluated to sufficient precision, we make the determination of acceptance or rejection. This is summarised in Algorithm 5, further noting that although the embedded loop is of undetermined length, by Corollary 3 we know that it halts in finite expected time ($\bar{K}$ can be interpreted as the expected computational cost of the nested loop, noting that $\mathbb{E}[\text{iterations}] := \sum_{i=0}^{\infty} i\mathbb{P}(\text{halt at step i}) = \sum_{i=0}^{\infty} \mathbb{P}(\text{halt at step i or later}) = \bar{K}$).

Algorithm 5 Simulating $W_q \sim \mathbb{W}_{s,\tau}^{W_s,W_\tau}|(W_s, W_\tau, \theta)$, given $q \in [s,\tau]$, the end points (W_s and the extremal value W_τ), and constrained such that $\forall u \in [s,\tau], W_u \in [W_s - \theta, W_s + \theta]$.

1. $\{b_q^{(i)}\}_{i=1}^3$: Simulate $b_q^{(1)}, b_q^{(2)}, b_q^{(3)} \overset{\text{iid}}{\sim} \text{N}\left(0, \frac{|\tau - q| \cdot |q - s|}{(\tau - s)^2}\right)$.
2. W_q: Set $W_q := W_\tau + (-1)^{\mathbb{1}(W_\tau < W_s)} \sqrt{(\tau - s)\left[\left(\frac{\theta(\tau - q)}{(\tau - s)^{3/2}} + b_q^{(1)}\right)^2 + (b_q^{(2)})^2 + (b_q^{(3)})^2\right]}$.
3. u: Simulate $u \sim \text{U}[0,1]$ and set $n = 1$.
4. $p^{\downarrow}, p^{\uparrow}$: While $u \notin [p_n^{\downarrow}, p_n^{\uparrow}]$, set $n = n + 1$.
5. p: If $u \leq p_n^{\downarrow}$ accept, else reject and return to Step 1.
6. Return (q, W_q).

C.3. Simulation of a single trajectory of constrained Brownian motion

We now have the constituent elements for Section 3, in which we simulate multivariate Brownian motion at any desired time marginal, with d-dimensional hypercuboids inscribing intervals of the state space in which the sample path almost surely lies (layers, more formally defined in Appendix B). Recall from Section 3 that the killing times are determined by a random variable whose distribution depends upon the inscribed layers, and so the presentation of Algorithm 6 necessitates a loop in which the determination of whether the stopping time occurs in the interval is required.

In the preceding subsections of Appendix C, we require the user-specified vector θ in order to determine the default hypercuboid inscription size. Therefore it is natural to set

$$\theta_j := \mathbf{\Lambda}_{j,j}^{1/2}\mathbf{c}. \tag{67}$$

Further note that due to the strong Markov property it is user preference as to whether this algorithm is run in its entirety for every required time marginal, or whether it resets layer information once one component breaches its boundary, re-initialises from that time on according to Algorithm 6 Step 4b.

Algorithm 6 Simulating constrained Brownian motion at a desired time marginal (t, W_t).

1. Input $\mathbf{W}_s$ and θ.
2. τ: For $i \in \{1, \dots, d\}$, simulate $(\tau^{(i)}, W_\tau^{(i)})$ as per Algorithm 4.
3. $\hat{\tau}$: Set $\hat{\tau} := \inf_i\{\tau^{(i)}\}$, set $j := \{i \in \{1, \dots, d\} : \tau^{(i)} = \hat{\tau}\}$.

* t: If required, simulate t as outlined in Section 3.

4. t: If $t \notin [s, \hat{\tau}]$,
   (a) $(\hat{\tau}, W_{\hat{\tau}}^{(\cdot)})$: For $i \in \{1, \dots, d\} \setminus j$, simulate $(\hat{\tau}, W_{\hat{\tau}}^{(i)})$ as per Algorithm 5.
   (b) $(\tau^{(j)}, W_\tau^{(j)})$: Simulate $(\tau^{(j)}, W_\tau^{(j)})$ as per Algorithm 4.
   (c) s: Set $s := \hat{\tau}$, and return to Step 3.
5. $(t, W_t^{(\cdot)})$: For $i \in \{1, \dots, d\}$, simulate $(t, W_t^{(i)})$ as per Algorithm 5.
6. Return (t, W_t).

Appendix D: Path-space Rejection Sampler (PRS) for $\mathbb{K}_T^{\mathbf{x}}$

In Section 3.1 we presented a path-space rejection sampler for $\mathbb{K}_T^{\mathbf{x}}$, in which a sample trajectory was drawn from Brownian motion measure, $\mathbf{X} \sim \mathbb{W}_T^{\mathbf{x}}$, and accepted with probability $P(\mathbf{X})$. Simulation of an event of probability $P(\mathbf{X})$ was achieved by simulating a number of binary events given by (22) and Theorem 2, with acceptance if all such events were accepted. The algorithmic pseudo-code for this approach is presented in Algorithm 7.

Algorithm 7 Path-space Rejection Sampler (PRS) for $\mathbb{K}_T^{\mathbf{x}}$ Algorithm

1. Input: $\mathbf{X}_0$.
2. R: Simulate layer information $R \sim \mathcal{R}$ as per Appendix C.
3. $P^{(1)}$: With probability $1 - \exp\{\Phi T - \sum_{i=1}^{n_R} L_{\mathbf{X}}^{(i)} \cdot [(\tau_i \wedge T) - \tau_{i-1}]\}$ reject and return to Step 2.
4. n_R: For i in $1 \to n_R$,
   (a) $\mathbb{U}_R^{(i)}$: Set $j = 0$, $\kappa_i = 0$, $\xi_0^{(i)} := \tau_{i-1}$ and $\zeta_1^{(i)} \sim \text{Exp}(U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)})$. While $\sum_j \zeta_j^{(i)} < [(\tau_i \wedge T) - \tau_{i-1}]$,
      i. $\xi_j^{(i)}$: Set $j = j + 1$ and $\xi_j^{(i)} = \xi_{j-1}^{(i)} + \zeta_j^{(i)}$.
      ii. $\mathbf{X}_{\xi_j^{(i)}}$: Simulate $\mathbf{X}_{\xi_j^{(i)}} \sim \text{MVN}(\mathbf{X}_{\xi_{j-1}^{(i)}}, (\xi_j^{(i)} - \xi_{j-1}^{(i)})\mathbf{\Lambda}) | R_{\mathbf{X}}^{(i)}$.
      iii. $P^{(2,i,j)}$: With probability $1 - [U_{\mathbf{X}}^{(i)} - \phi(\mathbf{X}_{\xi_j^{(i)}})]/[U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)}]$, reject path and return to Step 2.
      iv. $\zeta_{j+1}^{(i)}$: Simulate $\zeta_{j+1}^{(i)} \sim \text{Exp}(U_{\mathbf{X}}^{(i)} - L_{\mathbf{X}}^{(i)})$.
   (b) $\mathbf{X}_{\tau_i \wedge T}$: Simulate $\mathbf{X}_{\tau_i \wedge T} \sim \text{MVN}\,(\mathbf{X}_{\xi_j^{(i)}}, [(\tau_i \wedge T) - \xi_j^{(i)}]\mathbf{\Lambda}) | R_{\mathbf{X}}^{(i)}$.

Crucially, determination of acceptance is made using only a finite dimensional path skeleton. The PRS for $\mathbb{K}_T^{\mathbf{x}}$ outputs the skeleton composed of all intermediate simulations,

$$\mathcal{S}_{\mathrm{PRS}}(\mathbf{X}) := \left\{ \mathbf{X}_0, \left(\left(\xi_j^{(i)}, \mathbf{X}_{\xi_j^{(i)}} \right)_{j=1}^{\kappa_i}, R_{\mathbf{X}}^{(i)} \right)_{i=1}^{n_R} \right\}, \tag{68}$$

which is sufficient to simulate any finite-dimensional subset of the remainder of the sample path (denoted by $\mathbf{X}^{\mathrm{rem}}$) as desired without error (as outlined in [48, §3.1] and Appendix C),

$$\mathbf{X}_{(0,T)}^{\mathrm{rem}} \sim \otimes_{i=1}^{n_R} \left(\otimes_{j=1}^{\kappa_i} \mathbb{W}_{\xi_{j-1}^{(i)}, \xi_j^{(i)}}^{\mathbf{X}[\xi_{j-1}^{(i)}, \xi_j^{(i)}]} \right) \Big| R_{\mathbf{X}}^{(i)} \,. \tag{69}$$

Appendix E: Killed Brownian Motion (KBM)

In Section 3.1 we detailed an approach to simulate the killing time and location, $(\bar{\tau}, \mathbf{X}_{\bar{\tau}})$, for killed Brownian motion. To avoid unnecessary algorithmic complexity, note that we can recover the pair $(\bar{\tau}, \mathbf{X}_{\bar{\tau}})$ by a simple modification of Algorithm 7 in which we set $\forall i \, L_{\mathbf{X}}^{(i)} := \Phi$ and return the first rejection time. This is presented in Algorithm 8. A variant in which $L_{\mathbf{X}}^{(i)}$ is incorporated would achieve greater efficiency, but is omitted for notational clarity.

Algorithm 8 Killed Brownian Motion (KBM) Algorithm

1. Initialise: Set $i = 1$, $j = 0$, $\tau_0 = 0$. Input initial value $\mathbf{X}_0$.
2. R: Simulate layer information $R_{\mathbf{X}}^{(i)} \sim \mathcal{R}$ as per Appendix C, obtaining τ_i, $U_{\mathbf{X}}^{(i)}$.
3. ζ: Simulate $\zeta \sim \mathrm{Exp}(U_{\mathbf{X}}^{(i)} - \Phi)$.
4. ξ_j: Set $j = j + 1$ and $\xi_j = (\xi_{j-1} + \zeta) \wedge \tau_i$.
5. $\mathbf{X}_{\xi_j}$: Simulate $\mathbf{X}_{\xi_j} \sim \mathrm{MVN}(\mathbf{X}_{\xi_{j-1}}, (\xi_j - \xi_{j-1})\mathbf{\Lambda}) | R_{\mathbf{X}}^{(i)}$.
6. τ_i: If $\xi_j = \tau_i$, set $i = i + 1$ and return to Step 2.
7. P: With probability $[U_{\mathbf{X}}^{(i)} - \phi(\mathbf{X}_{\xi_j})]/[U_{\mathbf{X}}^{(i)} - \Phi]$ return to Step 3.
8. $(\bar{\tau}, \mathbf{X}_{\bar{\tau}})$: Return $(\bar{\tau}, \mathbf{X}_{\bar{\tau}}) = (\xi_j, \mathbf{X}_{\xi_j})$, $i_{\bar{\tau}} = i$, $j_{\bar{\tau}} = j$.

As in the PRS for $\mathbb{K}_T^{\mathbf{x}}$ presented in Appendix D, in KBM (Algorithm 8) we can recover in the interval $[0, \bar{\tau})$ the remainder of the sample path as desired without error as follows (where for clarity we have suppressed full notation, but can be conducted as described in Appendix C),

$$\mathcal{S}_{\mathrm{KBM}}(\mathbf{X}) := \left\{ \mathbf{X}_0, (\xi_j, \mathbf{X}_{\xi_j})_{j=1}^{j_{\bar{\tau}}}, (R_{\mathbf{X}}^{(i)})_{i=1}^{i_{\bar{\tau}}} \right\}, \qquad \mathbf{X}_{(0,T)}^{\mathrm{rem}} \sim \mathbb{W} | \mathcal{S}_{\mathrm{KBM}}. \tag{70}$$

Appendix F: Rejection Sampling based QSMC Algorithm

In Section 3.2 we considered the embedding of IS-KBM of Section 3.1 within SMC. A similar embedding for the rejection sampling variant (KBM) of Section 3.1 is considered here as the probability of the killed Brownian motion trajectory of Algorithm 8 remaining alive becomes arbitrarily small as diffusion time increases. As such, if one wanted to approximate the law of the process conditioned to remain alive until large T it would have prohibitive computational cost.

Considering the KBM algorithm presented in Appendix E, in which we simulate trajectories of killed Brownian motion, the most natural embedding of this within an SMC framework is to assign each particle constant un-normalised weight while alive, and zero weight when killed. Resampling in this framework simply consists of sampling killed particles uniformly at random from the remaining alive particle set. The manner in which we have constructed Algorithm 8 allows us to conduct this resampling in continuous time, and so we avoid the possibility of at any time having an alive particle set of size zero. We term this approach (Continuous Time) Rejection Quasi-Stationary Monte Carlo (R-QSMC), and present it in Algorithm 9. In Algorithm 9 we denote by $m(k)$ as a count of the number of killing events of particle trajectory k in the time elapsed until the m^{th} iteration of the algorithm.

Algorithm 9 (Continuous Time) Rejection Quasi-Stationary Monte Carlo Algorithm (R-QSMC) Algorithm.

1. **Initialisation Step ($m = 0$)**
   (a) Input: Starting value, $\hat{\mathbf{x}}$, number of particles, N.
   (b) $\mathbf{X}_0^{(\cdot)}$: For k in 1 to N set $\mathbf{X}_{t_0}^{(1:N)} = \hat{\mathbf{x}}$ and $w_{t_0}^{(1:N)} = 1/N$.
   (c) $\bar{\tau}_1^{(\cdot)}$: For k in 1 to N, simulate $\left(\bar{\tau}_1^{(k)}, \mathbf{X}_{\bar{\tau}_1}^{(k)}\right) \middle| \left(t_0^{(k)}, \mathbf{X}_{t_0}^{(k)}\right)$ as per Algorithm 8.
2. **Iterative Update Steps ($m = m + 1$)**
   (a) $\bar{\bar{\tau}}_m$: Set $\bar{\bar{\tau}}_m := \inf\{\{\bar{\tau}_{m(k)}^{(k)}\}_{k=1}^N\}$, $\bar{k} := \{k : \bar{\bar{\tau}}_m = \bar{\tau}_{m(k)}^{(k)}\}$.
   (b) K: Simulate $K \sim \mathrm{U}\{\{1, \ldots, n\} \setminus \bar{k}\}$.
   (c) $\mathbf{X}_{\bar{\bar{\tau}}_m}^{(\cdot)}$: Simulate $\mathbf{X}_{\bar{\bar{\tau}}_m}^{(\bar{k})} \sim \mathbb{W}|\mathcal{S}_{\mathrm{KBM}}^{(K)}$ as given by (70) and as per Algorithm 5.
   (d) $\bar{\bar{\tau}}_{m+1}$: Simulate $\left(\bar{\tau}_{m(\bar{k})+1}^{(\bar{k})}, \mathbf{X}_{\bar{\tau}_{m(\bar{k})+1}}^{(\bar{k})}\right) \middle| \left(\bar{\bar{\tau}}_m, \mathbf{X}_{\bar{\bar{\tau}}_m}^{(\bar{k})}\right)$ as per Algorithm 8.

Iterating the R-QSMC algorithm beyond some time t^* at which point we believe we have obtained convergence, and halting at time $T > t^*$, we can approximate the law of the killed process by the weighted occupation measures of the trajec-

tories (where $\forall t\, w_t^{(\cdot)} = 1/N$),

$$\pi(\mathrm{d}\mathbf{x}) \approx \hat{\pi}(\mathrm{d}\mathbf{x}) := \frac{1}{T-t^*}\int_{t^*}^{T}\sum_{k=1}^{N} w_t^{(k)} \cdot \delta_{\mathbf{X}_t^{(k)}}(\mathrm{d}\mathbf{x})\,\mathrm{d}t. \tag{71}$$

In some instances the tractable nature of Brownian motion will admit an explicit representation of (71). If not, one can simply sample the trajectories exactly at equally spaced points to find an unbiased approximation of (71), by means detailed in Appendix C.2 and Algorithm 4. In particular, if we let $t_0 := 0 < t_1 < \ldots < t_m := T$ such that $t_i - t_{i-1} := T/m$, then we can approximate the law of the killed process as we did in (29), where $w_{t*:T}^{(1:N)} = 1/N$.

Appendix G: Rejection sampling Scalable Langevin Exact (R-ScaLE) algorithm

In Section 4 we noted that the survival probability of a proposal Brownian motion sample path was related to the estimator $P(\mathbf{X})$ of Section 3.1, and in (32) of we develop a replacement estimator. The construction of control variates in Section 4.1 allows us to construct the replacement estimator such that it has good scalability properties. In a similar fashion to the embedding of this estimator within QSMC (Algorithm 2) resulting in ScaLE (Algorithm 3), we can embed this estimator with the rejection sampling variant R-QSMC (Algorithm 9) resulting in *Rejection Scalable Langevin Exact algorithm (R-ScaLE)* which we present in Algorithm 10.

Note as presented in Algorithm 10 we may also be concerned with the absolute growth of $\tilde{\Phi}$ (relative to Φ) as a function of n in order to study its computational complexity. Note however, as remarked upon in Appendix E, if this growth is not favourable one can modify Algorithm 8 to incorporate the additional path-space bound $\tilde{L}_{\mathbf{X}}^{(i)}$ for each layer. Details of this modification are omitted for notational clarity.

Algorithm 10 The R-ScaLE Algorithm (as per Algorithm 9 unless stated otherwise).

0. Choose $\hat{\mathbf{x}}$ and compute $\nabla \log \pi(\hat{\mathbf{x}})$, $\mathrm{div}\, \nabla \log \pi(\hat{\mathbf{x}})$. $\tilde{\Phi}$.

1c. On calling Algorithm 8

   (a) Replace Φ with $\tilde{\Phi}$.

   (b) Replace $U_{\mathbf{X}}^{(i)}$ in Step 2 with $\tilde{U}_{\mathbf{X}}^{(i)}$.

   (c) Replace Step 7 with: Simulate $I, J \overset{\text{iid}}{\sim} \mathrm{U}\{0,\ldots,n\}$, and with probability $[\tilde{U}_{\mathbf{X}}^{(i)} - \tilde{\phi}(\mathbf{X}_{\xi_j})]/[\tilde{U}_{\mathbf{X}}^{(i)} - \Phi]$ return to Step 3.

2d As Step 1c.

Appendix H: Discrete Time Sequential Monte Carlo Construction

Considering the discrete time system with state space $E_k = (C(h(k-1), hk], \mathcal{Z}_k)$ at discrete time k, with the process denoted $\mathfrak{X}_k = (X_{(h(k-1),hk]}, \mathfrak{Z}_k)$ in which the auxiliary variables $\mathfrak{Z}_k$ take values in some space $\mathcal{Z}_k$.

The ScaLE Algorithm, with resampling conducted deterministically at times $h, 2h, \ldots$ coincides exactly with the mean field particle approximation of a discrete time Feynman-Kac flow, in the sense and notation of [41], with transition kernel

$$M_k(\mathfrak{X}_{k-1}, d\mathfrak{X}_k) = \mathbb{W}^{X_{h(k-1)}}_{h(k-1),hk}(dX_{(h(k-1),hk]}) Q_k(X_{(h(k-1),hk]}, d\mathfrak{Z}_k)$$

and a potential function $G_k(\mathfrak{X}_k)$, which is left intentionally unspecified to allow a broad range of variants of the algorithm to be included, the property which it must possess to lead to a valid form of ScaLE Algorithm is specified below. Allowing

$$\overline{\mathbb{W}}^x_{0,hk}(\mathfrak{X}_{1:k}) = \mathbb{W}^{0,hk}_x(dX_{0:hk}) \prod_{i=1}^{k} Q_i(X_{(h(i-1),hi]}, d\mathfrak{Z}_i)$$

and specifying an extended version of the killed process via

$$\frac{d\overline{\mathbb{K}}^x_{0,hk}}{d\overline{\mathbb{W}}^x_{0,hk}}(\mathfrak{X}_{1:k}) \propto \prod_{i=1}^{k} G(\mathfrak{X}_i).$$

The validity of such a ScaLE Algorithm depends upon the following identity holding:

$$\frac{d\mathbb{K}^x_{0,hk}}{d\mathbb{W}^x_{0,hk}}(X_{0:hk}) \propto \mathbb{E}_{\mathbb{W}^x_{0,hk}}\left[\prod_{i=1}^{k} G_i(\mathfrak{X}_i) \middle| X_{0:hk} \right].$$

It is convenient to define some simplifying notation. We define the law of a discrete time process (in which is embedded a continuous time process taking values in $C[0,\infty)$):

$$\overline{\mathbb{W}}^x(d\mathfrak{X}) = \overline{\mathbb{W}}^x_{0,h}(d\mathfrak{X}_1) \prod_{k=1}^{\infty} \overline{\mathbb{W}}^{X_{h(k-1)}}_{h(k-1),hk}(d\mathfrak{X}_k)$$

and of a family of processes indexed by k, $\overline{\mathbb{K}}^x_k$, again incorporating a continuous time process taking values in $C[0,\infty)$, via:

$$\frac{d\overline{\mathbb{K}}^x_k}{d\overline{\mathbb{W}}_x}(\mathfrak{X}) \propto \prod_{i=1}^{k} G_i(\mathfrak{X}_i).$$

With a slight abuse of notation, we use the same symbol to refer to the associated finite dimensional distributions, with the intended distribution being indicated

by the argument. We also define the *marginal* laws, $\mathbb{W}^x$ and $\mathbb{K}_k^x$ via:

$$\begin{aligned}\mathbb{W}^x(dX) =& \overline{\mathbb{W}}^x(dX \times (\otimes_{p=1}^{\infty}\mathcal{Z}_p))\\ \mathbb{K}_k^x(dX) =& \overline{\mathbb{K}}_k^x(dX \times (\otimes_{p=1}^{\infty}\mathcal{Z}_p)).\end{aligned}$$

Proposition 2. *Under mild regularity conditions (cf. [41, 14]), for any $\varphi : \mathbb{R}^d \to \mathbb{R}$, any algorithm within the framework described admits a central limit in that:*

$$\lim_{N\to\infty}\sqrt{N}\left[\frac{1}{N}\sum_{i=1}^{N}\varphi(X_{hk}^i) - \mathbb{K}_k^x(\varphi(X_{hk}^i))\right] \Rightarrow \sigma_{k,G}(\varphi)Z$$

where, Z is a standard normal random variable, $\Rightarrow$ denotes convergence in distribution, and:

$$\begin{aligned}\sigma_k^2(\varphi) =& \mathbb{E}_{\overline{\mathbb{W}}}\left[\left(\frac{G_1(\mathfrak{X}_1)\mathbb{E}_{\overline{\mathbb{W}}^x}\left[\prod_{i=2}^{k}G(\mathfrak{X}_i)|\mathfrak{X}_1\right]}{\overline{\mathbb{W}}^x(\prod_{i=1}^{k}G(\mathfrak{X}_i))}\right)^2\mathbb{E}_{\mathbb{K}_k^x}\left[(\varphi(X_{hk}) - \mathbb{K}_k^x(\varphi(X_{hk})))^2\middle|X_h\right]\right] +\\ &\sum_{p=2}^{k-1}\mathbb{E}_{\overline{\mathbb{K}}_{p-1}^x}\left[\left(\frac{\overline{\mathbb{W}}^x(\prod_{i=0}^{p-1}G(\mathfrak{X}_i))}{\overline{\mathbb{W}}^x(\prod_{i=0}^{k}G(\mathfrak{X}_i))}G(\mathfrak{X}_p)\mathbb{E}_{\mathbb{W}^x}\left[\prod_{i=p+1}^{k}G(\mathfrak{X}_i)\middle|X_{hp}\right]\right)^2\mathbb{E}_{\mathbb{K}_k^x}\left[(\varphi(X_{hk}) - \mathbb{K}_k^x(\varphi(X_k)))^2\middle|X_{hp}\right]\right] +\\ &\mathbb{E}_{\overline{\mathbb{K}}_{k-1}^x}\left[\left(\frac{\overline{\mathbb{W}}^x(\prod_{i=0}^{k-1}G(\mathfrak{X}_i))}{\overline{\mathbb{W}}^x(\prod_{i=0}^{k}G(\mathfrak{X}_i))}G(\mathfrak{X}_k)\right)^2\left(\varphi(X_{hk}) - \overline{\mathbb{K}}_k^x(\varphi(X_{hk}))\right)^2\right]\end{aligned}$$

Proof Outline. It follows by a direct application of the argument underlying the Proposition of [28] (which itself follows from simple but lengthy algebraic manipulations from the results of [41, 14]) that for any test function, $\varphi : \mathbb{R}^d \to \mathbb{R}$ satisfying mild regularity conditions (cf. [41, 14]) that

$$\lim_{N\to\infty}\sqrt{N}\left[\frac{1}{N}\sum_{i=1}^{N}\varphi(X_{hk}^i) - \mathbb{K}_k^x(\varphi(X_{hk}^i))\right] \Rightarrow \sigma_{k,G}(\varphi)Z$$

where, Z is a standard normal random variable, $\Rightarrow$ denotes convergence in distribution, and:

$$\begin{aligned}\sigma_{k,G}^2(\varphi) =& \mathbb{E}_{\overline{\mathbb{W}}}\left[\left(\frac{d\overline{\mathbb{K}}_k^x}{d\overline{\mathbb{W}}^x}(X_{(0,h]},\mathfrak{Z}_1)\right)^2\mathbb{E}_{\overline{\mathbb{K}}_k^x}\left[\left(\varphi(X_{hk}) - \overline{\mathbb{K}}_k^x(\varphi(X_{hk}))\right)^2\middle|\overline{\mathcal{F}}_1\right]\right] +\\ &\sum_{p=2}^{k-1}\mathbb{E}_{\overline{\mathbb{K}}_{p-1}}\left[\left(\frac{d\overline{\mathbb{K}}_k^x}{d\overline{\mathbb{K}}_{p-1}^x}(X_{(0,hp]},\mathfrak{Z}_{1:p})\right)^2\mathbb{E}_{\overline{\mathbb{K}}_k^x}\left[\left(\varphi(X_{hk}) - \overline{\mathbb{K}}_k^x(\varphi(X_k))\right)^2\middle|\overline{\mathcal{F}}_p\right]\right] +\\ &\mathbb{E}_{\overline{\mathbb{K}}_{k-1}}\left[\left(\frac{d\overline{\mathbb{K}}_k^x}{d\overline{\mathbb{K}}_{k-1}}(X_{(0,hk]},\mathfrak{Z}_{1:k})\right)^2\left(\varphi(X_{hk}) - \overline{\mathbb{K}}_k^x(\varphi(X_{hk}))\right)^2\right]\end{aligned}$$

with $\{\overline{\mathcal{F}}_p\}_{p\geq 0}$ being the natural filtration associated with $\overline{\mathbb{W}}^x$.

This can be straightforwardly simplified to:

$$\begin{aligned}\sigma_k^2(\varphi) =& \mathbb{E}_{\overline{\mathbb{W}}}\left[\left(\frac{G_1(\mathfrak{X}_1)\mathbb{E}_{\overline{\mathbb{W}}^x}\left[\prod_{i=2}^k G(\mathfrak{X}_i)|\mathfrak{X}_1\right]}{\overline{\mathbb{W}}^x(\prod_{i=1}^k G(\mathfrak{X}_i))}\right)^2 \mathbb{E}_{\mathbb{K}_k^x}\left[\left(\varphi(X_{hk}) - \mathbb{K}_k^x(\varphi(X_{hk}))\right)^2\middle| X_h\right]\right] + \\ & \sum_{p=2}^{k-1}\mathbb{E}_{\overline{\mathbb{K}}_{p-1}^x}\left[\left(\frac{\overline{\mathbb{W}}^x(\prod_{i=0}^{p-1} G(\mathfrak{X}_i))}{\overline{\mathbb{W}}^x(\prod_{i=0}^k G(\mathfrak{X}_i))}G(\mathfrak{X}_p)\mathbb{E}_{\mathbb{W}^x}\left[\prod_{i=p+1}^k G(\mathfrak{X}_i)\middle| X_{hp}\right]\right)^2 \mathbb{E}_{\mathbb{K}_k^x}\left[\left(\varphi(X_{hk}) - \mathbb{K}_k^x(\varphi(X_k))\right)^2\middle| X_{hp}\right]\right] + \\ & \mathbb{E}_{\overline{\mathbb{K}}_{k-1}^x}\left[\left(\frac{\overline{\mathbb{W}}^x(\prod_{i=0}^{k-1} G(\mathfrak{X}_i))}{\overline{\mathbb{W}}^x(\prod_{i=0}^k G(\mathfrak{X}_i))}G(\mathfrak{X}_k)\right)^2 \left(\varphi(X_{hk}) - \overline{\mathbb{K}}_k^x(\varphi(X_{hk}))\right)^2\right]\end{aligned}$$

□

We conclude with the following corollary, showing that the particular combination of subsampling scheme and exact algorithm fits into this framework and providing its particular asymptotic variance expression.

Corollary 4. *Such a CLT is satisfied in particular:*

(a) If no subsampling is used and one evaluates the exact *(intractable) killing rate (as described in Algorithm 2).*
(b) If subsampling is employed within the construct of the layered path-space rejection sampler (as described in Algorithm 3).

Proof. Both claims follow directly by the above argument with the appropriate identifications.

(a) is established by setting:

$$\begin{aligned}\mathcal{Z}_k =& \emptyset \qquad\qquad G_k(\mathfrak{X}_k) = G(X_{[h(k-1),hk]}) \\ & \propto \frac{d\mathbb{K}_{h(k-1),hk}^{X_{h(k-1)}}}{d\mathbb{W}_{h(k-1),hk}^{X_{h(k-1)}}}(X_{(h(k-1):hk]})\end{aligned}$$

(b) is established by setting (where be denote by c the size of the subsampled

data):

$$\begin{aligned}
\mathcal{Z}_k =& \cup_{m_k=1}^{\infty} \otimes_{p=1}^{m_k} R(\tau_{k,p-1}, \tau_{k,p}) \\
R(s,t) =& \cup_{\kappa=0}^{\infty} \{\kappa\} \times (s,t]^{\kappa} \times \{1,\ldots,n\}^{2c\kappa} \\
\mathfrak{Z}_k =& (r_{k,1}, \ldots, r_{k,m_k}) \\
r_{k,p} =& (\kappa_{k,p}, \zeta_{k,p,1}, \ldots, \zeta_{k,p,\kappa_{k,p}}, s_{k,j,1,1:2c}, \ldots, s_{k,p,\kappa_{k,p},1:2c}) \\
G_k(\mathfrak{X}_k) =& \exp\left(-\sum_{p=1}^{m_k} L_\theta(X_{\tau_k,p-1})(\tau_{k,p} - \tau_{k,p-1})\right) \\
& \cdot \prod_{p=1}^{m_k} \prod_{j=1}^{\kappa_{k,p}} \left[\frac{U_\theta(X_{\tau_{k,p-1}}) - \widetilde{\phi}(X_{\zeta_{k,p,j}}, s_{k,p,j,1:2c})}{U_\theta(X_{\tau_{k,p-1}}) - L_\theta(X_{\tau_{k,p-1}})}\right] \\
Q_k(X_{(h(k-1),hk]}, d\mathfrak{Z}_k) =& \prod_{p=1}^{m_k} \left[\mathsf{PP}(d\zeta_{k,p,1:\kappa_{k,p}}; (U_\theta(X_{\tau_{k,p-1}}) - L_\theta(X_{\tau_{k,p-1}})), [\tau_{k,p-1}, \tau_{k,p}]) \right. \\
& \left. \cdot \prod_{j=1}^{\kappa_{k,p}} \frac{1}{n^{2c}} \prod_{l=1}^{2c} \delta_{\{1,\ldots,n\}}(ds_{k,j,l})) \right]
\end{aligned}$$

where $\mathsf{PP}(\cdot; \lambda, [a,b])$ denotes the law of a homogeneous Poisson process of rate λ over interval $[a,b]$, $\delta_{\{1,\ldots,n\}}$ denotes the counting measure over the first n natural numbers and a number of variables which correspond to deterministic transformations of the X process have been defined to lighten notation:

$$\tau_{k,p} = \begin{cases} (k-1)h & p = 0 \\ \inf\{t : |X_t - X_{\tau_{k,p-1}}| \geq \theta\} & p = 1, \ldots, m_k - 1 \\ kh & p = m_k \end{cases}$$

and m_k is the number of distinct layer pairs employed in interval k of the discrete time embedding of the algorithm (i.e. it is the number of first passage times simulated within the continuous time algorithm after time $(k-1)h$ until one of them exceeds kh; as detailed in Appendices C.1 and C.2). □